

Computationally Efficient Influence Maximization in Stochastic and Adversarial Models: Algorithms and Analysis

Justin Khim*

Varun Jog[†]Po-Ling Loh^{†‡}

November 2, 2016

Abstract

We consider the problem of influence maximization in fixed networks, for both stochastic and adversarial contagion models. The common goal is to select a subset of nodes of a specified size to infect so that the number of infected nodes at the conclusion of the epidemic is as large as possible. In the stochastic setting, the epidemic spreads according to a general triggering model, which includes the popular linear threshold and independent cascade models. We establish upper and lower bounds for the influence of an initial subset of nodes in the network, where the influence is defined as the expected number of infected nodes. Although the problem of exact influence computation is NP-hard in general, our bounds may be evaluated efficiently, leading to scalable algorithms for influence maximization with rigorous theoretical guarantees. In the adversarial spreading setting, an adversary is allowed to specify the edges through which contagion may spread, and the player chooses sets of nodes to infect in successive rounds. Both the adversary and player may behave stochastically, but we limit the adversary to strategies that are oblivious of the player's actions. We establish upper and lower bounds on the minimax pseudo-regret in both undirected and directed networks.

1 Introduction

Many data sets in contemporary scientific applications possess some form of network structure (Newman, 2003). Popular examples include data collected from social media websites such as Facebook and Twitter (Adamic and Adar, 2003; Liben-Nowell and Kleinberg, 2007), or electrical recordings gathered from a physical network of firing neurons (Sporns, 2011). In settings involving biological data, a common goal is to construct an abstract network representing interactions between genes, proteins, or other biomolecules (Hecker et al., 2009).

Over the last century, a vast body of work has been developed in the epidemiology literature to model the spread of disease (Kermack and McKendrick, 1927). The most popular models include SI (susceptible, infected), SIS (susceptible, infected, susceptible), and SIR (susceptible, infected, recovered), in which nodes may infect adjacent neighbors according to a certain stochastic process. These models have recently been applied to social network and viral marketing settings by computer scientists (Domingos and Richardson, 2001; Leskovec et al., 2007). In particular, the notion of influence, which refers to the expected number of infected individuals in a network at the conclusion of an epidemic spread, was studied by Kempe et al. (2003). However, determining an influence-maximizing seed set of a certain cardinality was shown to be NP-hard—in fact, even

*Department of Statistics, University of Pennsylvania, Philadelphia, PA 19104.

[†]Department of Electrical & Computer Engineering, University of Wisconsin, Madison, WI 53706.

[‡]Department of Statistics, University of Wisconsin, Madison, WI 53706.

computing the influence exactly in certain simple models is $\#P$ -hard (Chen et al., 2010a,b). Recent work in theoretical computer science has therefore focused on maximizing influence up to constant factors (Kempe et al., 2003, 2005; Borgs et al., 2014). In social networks, where edges represent interactions between individuals, the problem of influence maximization corresponds to identifying sets of individuals on which to impress an idea so that information spreads as widely as possible. Reliably modeling the transmission of information may also be relevant to the problems of predicting when knowledge becomes viral (e.g., in the case of online social networks), or limiting the spread of information through carefully positioned interventions (Cheng et al., 2014).

A series of recent papers (Lemonnier et al., 2014; Scaman et al., 2015; Lemonnier et al., 2016) establish computable upper bounds on the influence when information propagates in a stochastic manner according to an independent cascade model. In such a model, the infection spreads in rounds, and each newly infected node may infect any of its neighbors in the succeeding round. Central to these bounds is a matrix known as the hazard matrix, which encodes the transmission probabilities across edges in the graph. A recent paper by Lee et al. (2016) leverages “sensitive” edges in the network to obtain tighter bounds via a conditioning argument. The intention is that such bounds could be maximized to obtain a surrogate for the influence-maximizing set in the network. However, the tightness of the proposed bounds is yet unknown. The independent cascade model may be viewed as a special case of a more general triggering model, in which the infection status of each node in the network is determined by a random subset of neighbors (Kempe et al., 2003). The class of triggering models also includes another popular stochastic infection model known as the linear threshold model, and bounds for the influence function in linear threshold models have been explored in an independent line of work (Chen et al., 2010b; Zhou et al., 2014; Chen et al., 2009).

Naturally, one might wonder whether influence bounds might be derived for stochastic infection models in the broader class of triggering models, unifying and extending the aforementioned results. We answer this question affirmatively by establishing upper and lower bounds for the influence in general triggering models. Our derived bounds are attractive for two reasons: First, we are able to quantify the gap between our upper and lower bounds in the case of linear threshold models, expressed in terms of properties of the graph topology and edge probabilities governing the likelihood of infection. Second, maximizing a lower bound on the influence is guaranteed to yield a lower bound on the true maximum influence in the graph. Furthermore, as shown via the theory of submodular functions, the lower bounds in our paper may be maximized efficiently up to a constant-factor approximation via a greedy algorithm, leading to a highly-scalable algorithm with provable guarantees. To the best of our knowledge, the only previously established bounds for influence maximization are those mentioned above for the special cases of independent cascade and linear threshold models, and no theoretical or computational guarantees were known. This underscores the novelty of our contributions.

One shortcoming of the analysis of stochastic spreading models, however, is the fact that the edge weights characterizing the propensity for disease transmission are generally assumed to be known, which is unrealistic in many real-world applications. Several authors consider the “robust influence maximization” problem, where edge weights are only specified up to an interval or a set of small cardinality, and the goal is still to obtain a set of source vertices that approximately maximizes the true influence function (Chen et al., 2016a,b; He and Kempe, 2016). Chen et al. (2016a) also propose various sampling strategies applied to the edges of the graph that may be used to obtain increasingly accurate estimates for the true edge weights, where one is allowed to observe the transmission status for the edges that are sampled during multiple rounds of infection propagation.

To address this issue, we depart completely from the stochastic infection setting in the first half of our paper and analyze a different model in which the transmission (or lack of transmission) across edges in the network is dictated by an adversary. An infection is run in the graph in T separate

rounds, and the adversary is allowed to choose a different subset of edges in each round through which the disease will propagate. Note that the adversary may select edges in a randomized manner, so the class of strategies available to the adversary includes the stochastic spreading models described earlier; however, the general class of adversarial strategies studied in the second half of our paper is significantly larger. On each round, a player selects a set of source vertices, based only on knowledge of the subgraph of infected nodes in previous rounds. As in the stochastic spreading setting, the goal of the player is to infect as many vertices as possible, where vertices are considered to be infected if they either lie in the source set, or are reachable from the source set via a path of edges in the adversarial set. We define the regret of the player to be the difference between the total number of vertices infected by the best choice of source sets across rounds with full knowledge of the adversary’s actions and the total number of vertices infected by the player. These notions are taken from the theory of multi-armed bandits, but a key difference between the graph contagion setting and the standard multi-armed bandit setting is that in the latter case, the only information available to the player on each round is the reward obtained as a consequence of his or her actions. On the other hand, slightly more information is available to the player in our setting, since the player may be able to deduce additional information about which vertices would have been infected for a different choice of source vertices, based on observing the scope of the infection for a particular choice of source vertices (which reveals information about the edge set chosen by the adversary).

Our main contributions in this portion of the paper are to derive upper and lower bounds on the pseudo-regret for various adversarial and player strategies. Although the adversary and player are both allowed to adopt strategies involving randomization, we limit our attention to adversarial strategies that are oblivious to the successive actions of the player. We study both directed and undirected networks, where in the latter setting, contagion is allowed to spread in both directions when an edge is chosen by the adversary. In the case where the player only selects a single source vertex on each round, we derive upper bounds that hold uniformly over all possible adversarial strategies. Furthermore, we derive lower bounds for the minimax pseudo-regret when the underlying network is a complete graph, where the supremum is taken over all adversarial strategies and the infimum is taken over all player strategies. Our upper and lower bounds match up to constant factors in the case of directed networks. Notably, the bounds also agree with the usual rate for pseudo-regret in multi-armed bandits, showing that no new information is gained by the player by exploiting network structure. On the other hand, a gap exists between our upper and lower bounds for undirected networks, leaving open the possibility that the player may leverage the additional information from the network to incur less regret. Finally, we demonstrate how to extend our upper bounds to the setting where the player is allowed to choose multiple source vertices on each round. The proposed multi-source player strategy augments the source set sequentially using the single-source strategies as a subroutine, and is based on a general online greedy algorithm proposed by [Streeter and Golovin \(2007\)](#).

We now briefly mention other results in the literature addressing influence maximization in the framework of multi-armed bandits. Most work has focused on efficiently estimating edge weights in the stochastic infection setting by observing the extent of the infection in multiple runs with different choices of source vertices ([Chen et al., 2016b](#); [Vaswani et al., 2015](#); [Lei et al., 2015](#)). [Wen et al. \(2016\)](#) also study the problem of edge weight estimation in stochastic spreading models and discuss the notion of edge semi-bandit feedback, which we will adopt in our setting. Finally, [Narasimhan et al. \(2015\)](#) address the interesting question of accurately learning the influence function itself in a stochastic spreading model based on observing multiple rounds of infection. Importantly, none of these papers address a fully adversarial infection spreading setting, which is the topic of our investigation.

The remainder of our paper is organized as follows: In [Section 2](#), we rigorously characterize

the stochastic and adversarial spreading models to be analyzed in our paper. In Section 3, we present upper and lower bounds for the influence function in stochastic triggering models, along with approximation guarantees based on a sequential greedy algorithm. In Section 4, we present upper and lower bounds for pseudo-regret in the adversarial setting. Section 5 outlines the main ideas of the proofs, and Section 6 contains simulation results for stochastic influence maximization. We conclude the paper with a selection of open research questions in Section 7. Additional proof details may be found in the Appendix. Some of the results on stochastic spreading models are also contained in the conference version of this paper (Khim et al., 2016).

Notation. For a matrix $M \in \mathbb{R}^{n \times n}$, we write $\rho(M)$ to denote the spectral radius of M . We write $\|M\|_{\infty, \infty}$ to denote the ℓ_{∞} -operator norm of M . The matrix $\text{Diag}(M)$ denotes the matrix with diagonal entries equal to the diagonal entries of M and all other entries equal to 0. We write $\mathbf{1}_S$ to denote the all-ones vector supported on a set S . For a set A , let 2^A denote the power set of A . For the adversarial influence maximization setting, we will often use the notation \mathbb{E}_S to mean the expectation taken over the player’s actions for a fixed set of adversarial actions. This is the same as the conditional expectation with respect to the adversary’s actions. Similarly, we write \mathbb{E}_A to indicate the conditional expectation with respect to a fixed set of player actions.

2 Models of Infection

We first describe the infection spreading mechanisms we will study in this paper. In the first class of models, the disease spreads probabilistically across edges in discrete time steps. In the second class, an adversary decides (in a stochastic or deterministic manner) which edges in the graph will be open to contagion on each of T rounds, and the disease is transmitted across those edges according to contact with infected individuals. The goal is to determine which subsets of individuals to infect in an online manner in order to minimize a measure of “regret.”

2.1 Stochastic Models

We first introduce basic notation and define the stochastic infection models to be analyzed in our paper. The network of individuals is represented by a directed graph $G = (V, E)$, where V is the set of vertices and $E \subseteq V \times V$ is the set of edges. Furthermore, each directed edge (i, j) possesses a weight b_{ij} , whose interpretation changes with the specific model we consider. We denote the weighted adjacency matrix of G by $B = (b_{ij})$. We designate nodes as either *infected* or *uninfected* based on whether the information contagion has reached them. Let $\bar{A} := V \setminus A$.

2.1.1 Linear Threshold Models

We first describe the linear threshold model, introduced by Kempe et al. (2003). In this model, the edge weight b_{ij} denotes the influence node i has on node j . The probability that a node is infected depends on two quantities: the set of infected neighbors at a particular time instant, and a random node-specific threshold that remains constant over time. For each $i \in V$, we impose the condition $\sum_j b_{ji} \leq 1$.

The thresholds $\{\theta_i : i \in V\}$ are i.i.d. uniform random variables on $[0, 1]$. Beginning from an initially infected set $A \subseteq V$, the contagion proceeds in discrete time steps as follows: At every time step, each vertex i computes the total incoming weight from all infected neighbors, i.e., $\sum_{j \text{ is infected}} b_{ji}$. If this quantity exceeds θ_i , vertex i becomes infected. Once a node becomes infected, it remains infected for every succeeding time step. Note that the process necessarily stabilizes after

at most $|V|$ time steps. The expected size of the infection when the process stabilizes is known as the *influence* of A and is denoted by $\mathcal{I}(A)$. We may interpret the threshold θ_i as the level of immunity of node i .

Kempe et al. (2003) established the monotonicity and submodularity of the function $\mathcal{I} : 2^V \rightarrow \mathbb{R}$. As discussed in Section 3.4, these properties are key to the problem of influence maximization, which concerns maximizing $\mathcal{I}(A)$ for a fixed size of the set A . An important step used in describing the submodularity of \mathcal{I} is the “reachability via live-edge paths” interpretation of the linear threshold model. Since this interpretation is also crucial to our analysis, we describe it below.

Reachability via live-edge paths: Consider the weighted adjacency matrix B of the graph. We create a subgraph of G by selecting a subset of “live” edges, as follows: Each vertex i designates *at most one* incoming edge as a live edge, with edge (j, i) being selected with probability b_{ji} . (No neighboring edge is selected with probability $1 - \sum_j b_{ji}$.) The “reach” of a set A is defined as the set of all vertices i such that a path exists from A to i consisting only of live edges. The distribution of the set of nodes infected in the final state of the linear threshold model is identical to the distribution of reachable nodes under the live-edges model, when both are seeded with the same set A .

2.1.2 Independent Cascade Models

Kempe et al. (2003) also analyzed the problem of influence maximization in independent cascade models, a class of models motivated by interacting particle systems in probability theory (Durrett, 1988; Liggett, 2012). Similar to the linear threshold model, the independent cascade model begins with a set A of initially infected nodes in a directed graph $G = (V, E)$, and the infection spreads in discrete time steps. If a vertex i becomes infected at time t , it attempts to infect each uninfected neighbor j via the edge from i to j at time $t + 1$. The weight b_{ij} captures the probability that node i succeeds in infecting node j . This process continues until no more infections occur, which again happens after at most $|V|$ time steps.

The influence function for the independent cascade model was also shown to be monotonic and submodular, where the main step of the proof again relies on a “reachability via live-edge paths” interpretation. In this case, the description is straightforward: Given a graph G , every edge $(i, j) \in E$ is independently designated as a live edge with probability b_{ij} . It is then easy to see that the reach of A has the same distribution as the set of infected nodes in the final state of the independent cascade model.

2.1.3 Triggering Models

To unify the above models, Kempe et al. (2003) introduced the “triggering model,” which evolves as follows: Each vertex i chooses a random subset of its neighbors as *triggers*, where the choice of triggers for a given node is independent of the choice for all other nodes. If a node i is uninfected at time t but a vertex in its trigger set becomes infected, vertex i becomes infected at time $t + 1$. Note that the triggering model may be interpreted as a “reachability via live-edge paths” model if edge (j, i) is designated as live when i chooses j to be in its trigger set. The entry b_{ij} represents the probability that edge (i, j) is live. Clearly, the linear threshold and independent cascade models are special cases of the triggering model when the distributions of the trigger sets are chosen appropriately.

2.2 Adversarial Models

We now turn to a description of the adversarial spreading models to be studied in our paper. Let $G = (V, E)$ be a graph on n vertices, which may be directed or undirected. In the online problem,

we sequentially run T infection processes on G in which the player selects an infection source set \mathcal{S}_t with $|\mathcal{S}_t| = k$, for $t = 1, \dots, T$. For each t , the adversary designates a subset of edges $\mathcal{A}_t \subseteq E$ to be “open.” A node is considered to be infected at time t if and only if it is an element of \mathcal{S}_t or is reachable from \mathcal{S}_t via a path of open edges. Note that if G is an undirected graph, designating an edge as open allows an infection to spread in both directions. Furthermore, in the directed case, edges may exist in both directions between a given pair of nodes, in which case the adversary may designate both, one, or neither of the edges to be open. For an open edge set $A \subseteq E$ and infection source set $S \subseteq V$, we define $f(A, S)$ to be the fraction of vertices in the graph lying in the infected set.

The goal of the player is to devise a strategy that maximizes the aggregate number of infected nodes up to time T . Following standard bandit theory, we measure this in terms of the regret

$$R_T(\mathcal{A}, \mathcal{S}) = \sum_{t=1}^T f(\mathcal{A}_t, \mathcal{S}_*) - \sum_{t=1}^T f(\mathcal{A}_t, \mathcal{S}_t), \quad (1)$$

where

$$\mathcal{S}_* = \arg \max_{S: |S|=k} \sum_{t=1}^T f(\mathcal{A}_t, S)$$

is the optimal set that would be infected if the player were aware of the adversary’s strategy. (For additional background on multi-armed bandits, we refer the reader to [Cesa-Bianchi and Lugosi, 2006](#) and [Bubeck and Cesa-Bianchi, 2012](#) and the references cited therein.)

Since the player is unaware of the adversary’s strategy, however, it is often advantageous for the player to adopt a randomized strategy. Accordingly, we define the *expected regret*

$$\mathbb{E}[R_T(\mathcal{A}, \mathcal{S})] := \mathbb{E}_{\mathcal{A}, \mathcal{S}} \left[\max_{S: |S|=k} \sum_{t=1}^T f(\mathcal{A}_t, S) - \sum_{t=1}^T f(\mathcal{A}_t, \mathcal{S}_t) \right] \quad (2)$$

and *pseudo-regret*

$$\bar{R}_T(\mathcal{A}, \mathcal{S}) := \max_{S: |S|=k} \left\{ \mathbb{E}_{\mathcal{A}, \mathcal{S}} \left[\sum_{t=1}^T f(\mathcal{A}_t, S) - \sum_{t=1}^T f(\mathcal{A}_t, \mathcal{S}_t) \right] \right\}, \quad (3)$$

where the expectation in expressions (2) and (3) is taken with respect to potential randomization in both the strategies $\mathcal{A} = \{\mathcal{A}_t\}$ and $\mathcal{S} = \{\mathcal{S}_t\}$ of the adversary and player.

Next, we describe the classes of strategies for the adversary and player. The optimal strategy on each side depends on what information is available to the other. Throughout, we assume the adversary is *oblivious* of the player’s actions; i.e., at time $t = 0$, the adversary must decide on the (possibly random) strategy \mathcal{A} . Further note that if the adversary’s actions are deterministic rather than stochastic, we have the relationship

$$\bar{R}_T(\mathcal{A}, \mathcal{S}) = \mathbb{E}[R_T(\mathcal{A}, \mathcal{S})].$$

In general, however, we can only guarantee the bound

$$\bar{R}_T(\mathcal{A}, \mathcal{S}) \leq \mathbb{E}[R_T(\mathcal{A}, \mathcal{S})].$$

We use \mathcal{A} to denote the set of oblivious adversary strategies and \mathcal{A}_d to denote the set of deterministic adversary strategies.

Turning to the strategy of the player, we allow the player to choose his or her action at time t based on the feedback provided in response to the joint actions made by the player and adversary on preceding time steps. In particular, we assume that the player has knowledge of the connectivity of the underlying graph, and also receives *edge semi-bandit feedback* at each round (Wen et al., 2016). Rather than observing the entire set of open edges \mathcal{A}_t selected by the adversary at time t , this means the player observes the status of every edge (i, j) such that either i or j is in the reach of \mathcal{S}_t (in the undirected case), and the player observes the status of every edge (i, j) such that i is in the reach of \mathcal{S}_t (in the directed case). We use $\mathcal{J}(\mathcal{A}_t, \mathcal{S}_t)$ to denote the set of edges with status known to the player, and we also denote $\mathcal{J}^t = (\mathcal{J}(\mathcal{A}_1, \mathcal{S}_1), \dots, \mathcal{J}(\mathcal{A}_t, \mathcal{S}_t))$. Thus, a player strategy $\mathcal{S} = \{\mathcal{S}_t\}$ has the property that \mathcal{S}_t is a function of \mathcal{J}^{t-1} . We denote the set of all player strategies by \mathcal{P} , and we similarly denote the set of all deterministic player strategies by \mathcal{P}_d , meaning that \mathcal{S}_t is a deterministic function of \mathcal{J}^{t-1} . (Importantly, \mathcal{S}_t may be random, due to randomization in the adversary; however, *conditioned* on \mathcal{J}^{t-1} , the choice of \mathcal{S}_t is deterministic.)

Finally, we introduce the *scaled regret*

$$R^\alpha(\mathcal{A}, \mathcal{S}) = \alpha \sum_{t=1}^T f(\mathcal{A}_t, \mathcal{S}_*) - \sum_{t=1}^T f(\mathcal{A}_t, \mathcal{S}_t), \quad (4)$$

and the analogous quantities $\mathbb{E}[R^\alpha(\mathcal{A}, \mathcal{S})]$ and $\bar{R}_T^\alpha(\mathcal{A}, \mathcal{S})$. Our interest in the expression (4) is due to the fact that when $k > 1$, using a greedy algorithm for influence maximization is only guaranteed to achieve a $(1 - \frac{1}{e})$ -approximation of the truth. For this reason, we will examine the case when $\alpha = 1 - \frac{1}{e}$.

3 Bounds for Stochastic Models

We now state our results for influence maximization in triggering models. We begin with some notation, and then discuss upper and lower bounds for the influence function in linear threshold and general triggering models. We also establish submodularity properties of our lower bounds, which allow us to derive theoretical guarantees for optimization via a greedy algorithm.

3.1 Notation

Let B denote the weighted adjacency matrix of the graph. For a subset of vertices $A \subseteq V$, we write $b_{\bar{A}} \in \mathbb{R}^{|\bar{A}|}$ to denote the vector indexed by $i \in \bar{A}$, such that $b_{\bar{A}}(i) = \sum_{j \in A} b_{ji}$. Thus, $b_{\bar{A}}(i)$ records the total incoming weight from A into i . A *walk* in the graph G is a sequence of vertices $\{v_1, v_2, \dots, v_r\}$ such that $(v_i, v_{i+1}) \in E$, for $1 \leq i \leq r-1$. A *path* is a walk with no repeated vertices. If w_1 is a walk (or path) ending at a particular vertex, and w_2 is a walk (or path) starting from the same vertex, we write $w_1 w_2$ to denote the concatenation of w_1 and w_2 . For sets $S, T \subseteq V$, we write $\mathcal{W}_{S \rightarrow T}$ and $\mathcal{P}_{S \rightarrow T}$ to denote the set of all walks and paths, respectively, starting from a vertex in S and ending at a vertex in T . We define the *weight* of a walk to be $\omega(w) := \prod_{e \in w} b_e$, where the product is over all edges $e \in E$ included in w . (The weight of a walk of length 0 is defined to be 1.) For a set of walks $W = \{w_1, w_2, \dots, w_r\}$, we denote the sum of the weights of all walks in W by $\omega(W) = \sum_i \omega(w_i)$. The set of all walks (respectively, paths) from S to T lying entirely in a set U is denoted by $\mathcal{W}_{S \rightarrow T|U}$ (respectively, $\mathcal{P}_{S \rightarrow T|U}$). A superscript of k denotes paths or walks of length k ; e.g., $\mathcal{P}_{A \rightarrow B}^k$ denotes the set of all length- k paths from A to B . Note that $b_{\bar{A}}(i) = \omega(\mathcal{P}_{A \rightarrow i}^1)$.

3.2 Linear Threshold Models

We first provide upper and lower bounds for the influence of a set $A \subseteq V$ in the linear threshold model.

3.2.1 Upper Bound

We begin with upper bounds. We have the following result, proved in Section 5.1.1, which bounds the influence as a function of appropriate sub-blocks of the weighted adjacency matrix:

Theorem 1. *For any set $A \subseteq V$, we have the bound*

$$\mathcal{I}(A) \leq |A| + b_{\bar{A}}^T (I - B_{\bar{A}\bar{A}})^{-1} \mathbf{1}_{\bar{A}}. \quad (5)$$

In fact, the proof of Theorem 1 shows that the bound (5) may be strengthened to

$$\mathcal{I}(A) \leq |A| + b_{\bar{A}}^T \left(\sum_{i=1}^{n-|A|} B_{\bar{A},\bar{A}}^{i-1} \right) \mathbf{1}_{\bar{A}}, \quad (6)$$

since the upper bound is contained by considering paths from vertices in A to vertices in \bar{A} and summing over paths of various lengths (see also Theorem 4 below). It is also clear from the proof that the bound (6) is exact when the underlying graph G is a directed acyclic graph (DAG). However, the bound (5) may be preferable in some cases from the point of view of computation or interpretation.

3.2.2 Lower Bounds

We also establish lower bounds on the influence. The following theorem, proved in Section 5.1.2, provides a family of lower bounds, indexed by $m \geq 1$:

Theorem 2. *For any $m \geq 1$, we have the following lower bound on the influence of A :*

$$\mathcal{I}(A) \geq \sum_{k=0}^m \omega(P_A^k), \quad (7)$$

where P_A^k is the set of all paths from A to \bar{A} of length k , such that only the starting vertex lies in A . We note some special cases when the bounds may be written explicitly:

$$m = 1 : \quad \mathcal{I}(A) \geq |A| + b_{\bar{A}}^T \mathbf{1}_{\bar{A}} := LB_1(A) \quad (8)$$

$$m = 2 : \quad \mathcal{I}(A) \geq |A| + b_{\bar{A}}^T (I + B_{\bar{A},\bar{A}}) \mathbf{1}_{\bar{A}} := LB_2(A) \quad (9)$$

$$m = 3 : \quad \mathcal{I}(A) \geq |A| + b_{\bar{A}}^T (I + B_{\bar{A},\bar{A}} + B_{\bar{A},\bar{A}}^2 - \text{Diag}(B_{\bar{A},\bar{A}}^2)) \mathbf{1}_{\bar{A}}. \quad (10)$$

Remark: As noted in Chen et al. (2010b), computing the exact influence function is $\#P$ -hard precisely because it is difficult to write down an expression for $\omega(P_A^k)$ for arbitrary values of k . When $m > 3$, we may use the techniques in Movarraei and Shikare (2014) and Movarraei and Boxwala (2015a,b) to obtain explicit lower bounds when $m \leq 7$. Note that as m increases, the sequence of lower bounds approaches the true value of $\mathcal{I}(A)$.

The lower bound (8) has a very simple interpretation: When $|A|$ is fixed, the function $LB_1(A)$ computes the aggregate weight of edges from A to \bar{A} . Furthermore, as established in Theorem 7 below, the function LB_1 is monotonic. Hence, maximizing LB_1 with respect to A is equivalent to finding a maximum cut in the directed graph. (For more details, see Section 3.4.) The lower bounds (9) and (10) also take into account the weight of paths of length 2 and 3 from A to \bar{A} .

3.2.3 Closeness of Bounds

A natural question concerns the proximity of the upper bound (5) to the lower bounds in Theorem 2. The bounds may be far apart in general, as illustrated by the following example:

Example: Consider a graph G with vertex set $\{1, 2, \dots, n\}$, and edge weights given by

$$w_{ij} = \begin{cases} 0.5, & \text{if } i = 1 \text{ and } j = 2, \\ 0.5, & \text{if } i = 2 \text{ and } 3 \leq j \leq n, \\ 0, & \text{otherwise.} \end{cases}$$

Let $A = \{1\}$. We may check that $LB_1(A) = 1.5$. Furthermore, $\mathcal{I}(A) = \frac{n+2}{4}$, and any upper bound necessarily exceeds this quantity. Hence, the gap between the upper and lower bounds may grow linearly in the number of vertices. (Similar examples may be computed for LB_2 , as well.)

The reason for the linear gap in the above example is that vertex 2 has a very large *outgoing* weight; i.e., it is highly infectious. Our next result shows that if the graph does not contain any highly-infectious vertices, the upper and lower bounds are guaranteed to differ by a constant factor. The result is stated in terms of the maximum row sum $\lambda_{\bar{A},\infty} = \|B_{\bar{A},\bar{A}}\|_{\infty,\infty}$, which corresponds to the maximum outgoing weight of the nodes in \bar{A} . The proof is contained in Section 5.1.3.

Theorem 3. *Suppose $\lambda_{\bar{A},\infty} < 1$. Then*

$$\frac{UB}{LB_1} \leq \frac{1}{1 - \lambda_{\bar{A},\infty}}, \quad \text{and} \quad \frac{UB}{LB_2} \leq \frac{1}{1 - \lambda_{\bar{A},\infty}^2}.$$

Since the column sums of B are bounded above by 1 in a linear threshold model, we have the following corollary:

Corollary 1. *Suppose B is symmetric and $A \subsetneq V$. Then*

$$\frac{UB}{LB_1} \leq \frac{1}{1 - \lambda_{\bar{A},\infty}}, \quad \text{and} \quad \frac{UB}{LB_2} \leq \frac{1}{1 - \lambda_{\bar{A},\infty}^2}.$$

Note that if $\lambda_\infty = \|B\|_{\infty,\infty}$, we certainly have $\lambda_{\bar{A},\infty} \leq \lambda_\infty$ for any choice of $A \subseteq V$. Hence, Theorem 3 and Corollary 1 hold *a fortiori* with $\lambda_{\bar{A},\infty}$ replaced by λ_∞ .

3.3 Triggering Models

We now generalize our discussion to the broader class of triggering models. Recall that in this model, b_{ij} records the probability that (i, j) is a live edge.

3.3.1 Upper Bound

We begin by deriving an upper bound, which shows that inequality (6) holds for any triggering model:

Theorem 4. *In a general triggering model, the influence of $A \subseteq V$ satisfies inequality (6).*

The proof of Theorem 4 is contained in Section 5.2.1.

Note that the approach we use for general triggering models relies on slightly more sophisticated observations than the proof for linear threshold models. Furthermore, the finite sum in inequality (6) may not in general be replaced by an infinite sum, as in the statement of Theorem 1 for the case of linear threshold models. This is because if $\rho(B_{\bar{A},\bar{A}}) > 1$, the infinite series will not converge.

3.3.2 Lower Bound

We now establish a general lower bound. The following theorem is proved in Section 5.2.2:

Theorem 5. *Let $A \subseteq V$. The influence of A satisfies the inequality*

$$\mathcal{I}(A) \geq \sum_{i \in V} \sup_{p \in P_{A \rightarrow i}} \omega(p) := LB_{\text{trig}}(A), \quad (11)$$

where $P_{A \rightarrow i}$ is the set of all paths from A to i such that only the starting vertex lies in A .

The proof of Theorem 5 shows that the bound (11) is sharp when at most one path exists from A to each vertex i . In the case of linear threshold models, the bound (11) is not directly comparable to the bounds stated in Theorem 2, since it involves maximal-weight paths rather than paths of certain lengths. Hence, situations exist in which one bound is tighter than the other, and vice versa (e.g., see the example in Section 3.2.3).

3.3.3 Independent Cascade Models

We now apply the general bounds obtained for triggering models to the case of independent cascade models. Theorem 4 implies the following “worst-case” upper bounds on influence, which only depend on $|A|$. The proof of the theorem is contained in Section 5.2.3:

Theorem 6. *The influence of $A \subseteq V$ in an independent cascade model satisfies*

$$\mathcal{I}(A) \leq |A| + \lambda_\infty |A| \cdot \frac{1 - \lambda_\infty^{n-|A|}}{1 - \lambda_\infty}. \quad (12)$$

In particular, if $\lambda_\infty < 1$, we have

$$\mathcal{I}(A) \leq \frac{|A|}{1 - \lambda_\infty}. \quad (13)$$

Note that when $\lambda_\infty > 1$, the bound (12) exceeds n for all large enough n , so the bound is trivial.

It is instructive to compare Theorem 6 with the results of Lemonnier et al. (2016). The *hazard matrix* of an independent cascade model with weighted adjacency matrix (b_{ij}) is defined by

$$\mathcal{H}_{ij} = -\log(1 - b_{ij}), \quad \forall (i, j).$$

The following result is stated in terms of the spectral radius $\rho = \rho\left(\frac{\mathcal{H} + \mathcal{H}^T}{2}\right)$:

Proposition 1 (Corollary 1 in Lemonnier et al., 2016). *Let $A \subsetneq V$, and suppose $\rho < 1 - \delta$, where $\delta = \left(\frac{|A|}{4(n-|A|)}\right)^{1/3}$. Then $\mathcal{I}(A) \leq |A| + \sqrt{\frac{\rho}{1-\rho}} \sqrt{|A|(n-|A|)}$.*

As illustrated in the following example, the bound in Theorem 6 may be significantly tighter than the bound provided in Proposition 1:

Example: Consider a directed Erdős-Rényi graph on n vertices, where each edge (i, j) is independently present with probability $\frac{c}{n}$. Suppose $c < 1$. For any set $|A|$, the bound (13) gives

$$\mathcal{I}(A) \leq \frac{|A|}{1 - c}. \quad (14)$$

It is easy to check that $\rho\left(\frac{\mathcal{H}+\mathcal{H}^T}{2}\right) = -(n-1)\log\left(1 - \frac{c}{n}\right)$. For large values of n , we have $\rho(\mathcal{H}) \rightarrow c < 1$, so Proposition 1 implies the (approximate) bound

$$\mathcal{I}(A) \leq |A| + \sqrt{\frac{c}{1-c}} \sqrt{|A|(n-|A|)}.$$

In particular, this bound increases with n , unlike our bound (14). Although the example is specific to Erdős-Rényi graphs, we conjecture that whenever $\|B\|_{\infty,\infty} < 1$, the bound in Theorem 6 is tighter than the bound in Proposition 1.

3.4 Maximizing Influence

We now turn to the question of choosing a set $A \subseteq V$ of cardinality at most k that maximizes $\mathcal{I}(A)$. We begin by reviewing the notion of submodularity, which will be crucial in our discussion of influence maximization algorithms. We have the following definition:

Definition 1 (Submodularity). *A set function $f : 2^V \rightarrow \mathbb{R}$ is submodular if either of the following equivalent conditions holds:*

(i) *For any two sets $S, T \subseteq V$,*

$$f(S \cup T) + f(S \cap T) \leq f(S) + f(T). \quad (15)$$

(ii) *For any two sets $S \subseteq T \subseteq V$ and any $x \notin T$, the following inequality holds:*

$$f(T \cup \{x\}) - f(T) \leq f(S \cup \{x\}) - f(x). \quad (16)$$

The left and right sides of inequality (16) are the discrete derivatives of f evaluated at T and S .

Submodular functions arise in a wide variety of applications. Although submodular functions resemble convex and concave functions, optimization may be quite challenging; in fact, many submodular function maximization problems are NP-hard. However, positive submodular functions may be maximized efficiently if they are also monotonic, where monotonicity is defined as follows:

Definition 2 (Monotonicity). *A function $f : 2^V \rightarrow \mathbb{R}$ is monotonic if for any two sets $S \subseteq T \subseteq V$,*

$$f(S) \leq f(T).$$

Equivalently, a function is monotonic if its discrete derivative is nonnegative at all points.

We have the following celebrated result, which guarantees that the output of the greedy algorithm provides a $(1 - \frac{1}{e})$ -approximation to the cardinality-constrained maximization problem:

Proposition 2 (Theorem 4.2 of Nemhauser and Wolsey, 1978). *Let $f : 2^V \rightarrow \mathbb{R}_+$ be a monotonic submodular function. For any $k \geq 0$, define $m^*(k) = \max_{|S| \leq k} f(S)$. Suppose we construct a sequence of sets $\{S_0 = \phi, S_1, \dots, S_k\}$ in a greedy fashion, such that $S_{i+1} = S_i \cup \{x\}$, where x maximizes the discrete derivative of f evaluated at S_i . Then*

$$f(S_k) \geq \left(1 - \frac{1}{e}\right) m^*(k).$$

3.4.1 Greedy Algorithms

Kempe et al. (2003) leverage Proposition 2 and the submodularity of the influence function to derive guarantees for a greedy algorithm for influence maximization in the linear threshold model. However, due to the intractability of exact influence calculations, each step of the greedy algorithm requires approximating the influence of several augmented sets. This leads to an overall runtime of $\mathcal{O}(nk)$ times the runtime for simulations, and introduces an additional source of error.

As the results of this section establish, the lower bounds $\{LB_m\}_{m \geq 1}$ and LB_{trig} appearing in Theorems 2 and 5 are also conveniently submodular, implying that Proposition 2 applies. In contrast to the algorithm studied by Kempe et al. (2003), however, our proposed greedy algorithms do not involve expensive simulations, since LB_m (for small values of m) and LB_{trig} are straightforward to evaluate exactly. This means the resulting algorithm is extremely fast to compute even on large networks. The following theorems are proved in Appendices A.2.1 and A.2.2, respectively:

Theorem 7. *The functions $\{LB_m\}_{m \geq 1}$ are monotone and submodular. Thus, for any $k \leq n$, a greedy algorithm that maximizes LB_m at each step yields a $(1 - \frac{1}{e})$ -approximation to $\max_{A \subseteq V: |A| \leq k} LB_m(A)$.*

Theorem 8. *The function LB_{trig} is monotone and submodular. Thus, for any $k \leq n$, a greedy algorithm that maximizes LB_{trig} at each step yields a $(1 - \frac{1}{e})$ -approximation to $\max_{A \subseteq V: |A| \leq k} LB_{\text{trig}}(A)$.*

Note that maximizing $LB_m(A)$ or $LB_{\text{trig}}(A)$ necessarily provides a lower bound on $\max_{A \subseteq V} \mathcal{I}(A)$.

3.4.2 Convex Relaxations

Greedy algorithms only provide one approach to tackling combinatorial optimization problems; it is natural to wonder whether a non-greedy approach could furnish even better empirical performance and/or theoretical guarantees. Indeed, the optimization problem $\max_{A \subseteq V: |A| \leq k} LB_1$ leads to several natural convex relaxations. We briefly discuss some such relaxations for maximizing LB_1 subject to a cardinality constraint.

By the monotonicity of LB_1 (cf. Lemma 2), we know that maximizing $LB_1(A)$ over the set $\{A : |A| \leq k\}$ is equivalent to the following optimization problem:

$$\begin{aligned} \max \quad & LB_1(A) \\ \text{s.t.} \quad & |A| = k. \end{aligned} \tag{17}$$

As remarked in Section 3.2.2, the program (17) is exactly equivalent to the maximum directed cut with given sizes of parts (MAX DICUT WITH GSP) problem, which has been studied extensively in the theoretical computer science literature.

In particular, the theoretical guarantees corresponding to various approximation algorithms for MAX DICUT WITH GSP may be applied directly to the problem of maximizing LB_1 . We mention two state-of-the-art approaches here:

1. Ageev et al. (2001) provide an algorithm based on an LP relaxation of an integer programming formulation of the optimization problem. Since the solutions may be non-integral, the resulting vectors are post-processed via a method known as pipage rounding. The final overall procedure is guaranteed to produce a $\frac{1}{2}$ -approximation of the true maximum.
2. Gupta et al. (2010) propose a method based on a type of modified greedy procedure. It employs an approximation algorithm by Feige et al. (2011) as a subroutine. Elements are again

added sequentially to the active set, but a choice may be made among several options for the augmented element. The overall algorithm is guaranteed to return a $\frac{1}{4+\alpha}$ -approximation for the true maximum.

Of course, one algorithm may perform better than another in practice even if it possesses weaker theoretical guarantees. We leave a proper investigation of the empirical performance of the above algorithms, in comparison to the performance of the greedy algorithm discussed in Section 3.4.1, for future work.

Finally, we remark that in the case when incoming weighted edges at a particular node may sum to more than one (i.e., the row sums of B exceed 1), the objective function $LB_1(A)$ in the optimization problem (17) is no longer guaranteed to be submodular or monotone. Thus, none of the aforementioned results are applicable. However, the program (17) still admits a convenient SDP relaxation, analogous to the SDP relaxation derived by [Goemans and Williamson \(1995\)](#) for the MAX DICUT problem. To the best of our knowledge, theoretical guarantees for SDP relaxations of the constrained MAX DICUT problem are yet unknown.

4 Bounds for Adversarial Models

We now shift our attention to adversarial spreading models. Recall that our goal in this setting is to control the (scaled) pseudo-regret, defined by

$$\overline{R}_T^\alpha(\mathcal{A}, \mathcal{S}) = \min_{S: |S|=k} \left\{ \mathbb{E}_{\mathcal{A}, \mathcal{S}} \left[\alpha \sum_{t=1}^T f(\mathcal{A}_t, S) - \sum_{t=1}^T f(\mathcal{A}_t, \mathcal{S}_t) \right] \right\},$$

where $\alpha = 1$ for the unscaled version. In this section, we provide upper and lower bounds for the pseudo-regret. Specifically, we focus on the quantity

$$\inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}} \overline{R}_T^\alpha(\mathcal{A}, \mathcal{S}),$$

where the supremum is over the class of strategies of an oblivious adversary, and the infimum is over the class of strategies for a player with edge semi-bandit feedback. Note that if we instead take a supremum over deterministic adversary strategies, we have the relation

$$\inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}_d} \overline{R}_T^\alpha(\mathcal{A}, \mathcal{S}) = \inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}_d} \mathbb{E}[R_T^\alpha(\mathcal{A}, \mathcal{S})],$$

since the pseudo-regret agrees with the expected regret; however, in the more general case, we may only deduce lower (but not upper) bounds on the minimax expected regret from bounds on the pseudo-regret.

A rough outline of our approach is as follows: We establish upper bounds by presenting particular strategies for the player that ensure an appropriately bounded regret under all adversarial strategies. For lower bounds, the general technique is to provide an ensemble of possible actions for the adversary that are difficult for the player to distinguish in the influence maximization problem, which forces the player to incur a certain level of regret.

4.1 Undirected Graphs

We begin by deriving regret upper bounds for undirected graphs. We initially restrict our attention to the case $k = 1$. The proposed player strategy for $k > 1$, and corresponding regret bounds, will build upon the results in the single-source setting.

4.1.1 Upper Bounds for a Single Source

Consider a randomized player strategy that selects $\mathcal{S}_t = \{i\}$ with probability $p_{i,t}$. We first describe the Exp3 algorithm, which computes the probabilities $p_{i,t}$ in an online manner based on a running estimate of the loss incurred at node i up to time t :

Exp3 with loss estimates $\{\widehat{\ell}_{i,t}\}$

Given: A nonincreasing sequence of real numbers $\{\eta_t\}_{t=1}^T$.

Output: A stochastic player strategy $\{\mathcal{S}_t\}$.

Let p_1 be the uniform distribution over $\{1, \dots, n\}$.

For each round $t = 1, \dots, T$:

- (1) Draw a vertex \mathcal{S}_t from the distribution p_t .
- (2) For each vertex $i = 1, \dots, n$, compute the estimated cumulative loss

$$\widehat{L}_{i,t} = \widehat{L}_{i,t-1} + \widehat{\ell}_{i,t}.$$

- (3) Compute the new distribution $p_{t+1} = (p_{1,t+1}, \dots, p_{n,t+1})$ over the vertices, where

$$p_{i,t+1} = \frac{\exp(-\eta_t \widehat{L}_{i,t})}{\sum_{k=1}^n \exp(-\eta_t \widehat{L}_{k,t})}.$$

Note that Exp3 depends on the choice of loss estimates $\{\widehat{\ell}_{i,t}\}$ and parameters $\{\eta_t\}$. The most basic loss estimate, which follows from standard bandit theory and ignores all information about the graph, is

$$\widehat{\ell}_{i,t}^{\text{node}} = \frac{\ell_{i,t}}{p_{i,t}} \mathbf{1}_{\mathcal{S}_t = \{i\}}, \quad (18)$$

where $\ell_{i,t} = 1 - f(\mathcal{A}_t, \{i\})$ is the loss incurred if the player were to choose $\mathcal{S}_t = \{i\}$. Importantly, $\widehat{\ell}_{i,t}^{\text{node}}$ is always computable for any choice the player makes at time t , and is an unbiased estimate of $\ell_{i,t}$.

On the other hand, if $\mathcal{S}_t = \{i\}$ and another node j is infected (i.e., in the connected component formed by the open edges of \mathcal{A}_t), the player also knows the loss that would have been incurred if $\mathcal{S}_t = \{j\}$, since $f(\mathcal{A}_t, \{i\}) = f(\mathcal{A}_t, \{j\})$. This motivates an alternative loss estimate that is nonzero even when $\mathcal{S}_t \neq \{i\}$. In particular, we may express

$$\ell_{i,t} = \frac{1}{n} \sum_{j \neq i} \ell_{i,j}^t, \quad (19)$$

where $\ell_{i,j}^t = 1$ is the indicator that i and j are in different connected components formed by the open edges of \mathcal{A}_t . We then define

$$\widehat{\ell}_{i,t}^{\text{sym}} = \frac{1}{n} \sum_{j \neq i} \ell_{i,j}^t \frac{Z_{ij}}{p_{i,t} + p_{j,t}},$$

where $Z_{ij} = \mathbf{1}_{\mathcal{S}_t \cap \{i,j\} \neq \emptyset}$. Using equation (19), it is easy to see that $\hat{\ell}_{i,t}^{\text{sym}}$ is an unbiased estimate for $\ell_{i,t}$. Furthermore, the estimator $\hat{\ell}_{i,t}^{\text{sym}}$ is always computable by the player, since the value of $\ell_{i,j}^t$ is known by the player whenever \mathcal{S}_t is known. We call $\hat{\ell}_{i,t}^{\text{sym}}$ the *symmetric loss*. We then have the following theorem:

Theorem 9 (Symmetric loss, Exp3). *Suppose $n > 1$, and the player uses the strategy $\mathcal{S}_{\text{Exp3}}^{\text{sym}}$ corresponding to Exp3 with the symmetric loss $\hat{\ell}^{\text{sym}}$ and $\eta_t = \sqrt{\frac{4 \log n}{T(n+1)}}$. Then the pseudo-regret satisfies the bound*

$$\sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}_{\text{Exp3}}^{\text{sym}}) \leq \sqrt{T(n+1) \log n}.$$

Remark 1. *It is instructive to compare the result of Theorem 9 with analogous regret bounds for generic multi-armed bandits. When the Exp3 algorithm is run with the loss estimates (18), standard analysis (Bubeck and Cesa-Bianchi, 2012) establishes an upper bound of $\sqrt{2Tn \log n}$. Thus, using the symmetric loss, which leverages the graphical nature of the problem, produces slight gains.*

In fact, more sophisticated player strategies lead to tighter upper bounds on the pseudo-regret. Bubeck and Cesa-Bianchi (2012) suggest a method based on the Online Stochastic Mirror Descent (OSMD) algorithm, which is specified by loss estimates $\{\ell_{i,t}\}$ and learning rates $\{\eta_t\}$, as well as a Legendre function F . The Exp3 algorithm described above may be viewed as a special case of OSMD with particular choice of Legendre function. In order to avoid excessive technicalities, we defer the details of the OSMD algorithm to Section 5.3.2, and for now, state the following regret bounds:

Theorem 10 (Symmetric loss, OSMD). *Suppose the player uses the strategy $\mathcal{S}_{\text{OSMD}}^{\text{sym}}$ corresponding to OSMD with the symmetric loss $\hat{\ell}^{\text{sym}}$ and appropriate parameters. Then the pseudo-regret satisfies the bound*

$$\sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}_{\text{OSMD}}^{\text{sym}}) \leq 2^{\frac{1}{4}} \sqrt{Tn}.$$

Remark 2. *The chief difference between the bounds based on the OSMD and Exp3 algorithms is the removal of the $\sqrt{\log n}$ term from Theorems 9 to 10. Again, such a phenomenon is well-known in the bandit literature as well, and running the OSMD algorithm with the loss estimates (18) would result in a pseudo-regret upper bound of $2^{\frac{3}{2}} \sqrt{Tn}$.*

4.1.2 Upper Bounds for Multiple Sources

We now turn to the case $k > 1$, where the player chooses multiple source vertices at each time step. As discussed in Section 2.2, we are interested in bounding the scaled pseudo-regret $\bar{R}_T^\alpha(\mathcal{A}, \mathcal{S})$ with $\alpha = 1 - \frac{1}{e}$, since it is difficult to maximize the influence even in an offline setting, and the greedy algorithm is only guaranteed to provide a $(1 - \frac{1}{e})$ -approximation of the truth.

Our proposed player strategy is based on an online greedy adaptation of the strategy used in the single-source setting. We assume the player is allowed to choose source vertices sequentially at time t and observes edge semi-bandit feedback immediately after each selection. The algorithm, inspired by Streeter and Golovin (2007), is outlined below:

Online Greedy Algorithm

Given: A single-source player strategy \mathcal{S}^1 .

Output: A k -source player strategy $\mathcal{S}^k = \{\mathcal{S}_t\}_{1 \leq t \leq T}$.

For each $t = 1, \dots, T$, choose $\mathcal{S}_t = \{v_{1,t}, \dots, v_{k,t}\}$ sequentially, as follows:

- (1) Select $v_{1,t}$ according to the single-source strategy \mathcal{S}^1 .
- (2) For each $i > 1$, select $v_{i,t}$ according to the single-source strategy \mathcal{S}^1 , based on the edge semi-bandit feedback $\mathcal{J}(\mathcal{A}_t, \{v_{1,t}, \dots, v_{i,t}\}) \setminus \mathcal{J}(\mathcal{A}_t, \{v_{1,t}, \dots, v_{i-1,t}\})$.

In other words, the Online Greedy Algorithm runs the player's strategy for single-source selection k times in parallel, with losses computed marginally for each successively chosen vertex. The “greedy” component of the algorithm corresponds to the fact that the player makes a selection of the set of i^{th} source vertices in the best possible way based on the information available (i.e., according to the single-source strategy that is designed to incur a small pseudo-regret). Note that the feedback

$$\mathcal{J}(\mathcal{A}_t, \{v_{1,t}, \dots, v_{i,t}\}) \setminus \mathcal{J}(\mathcal{A}_t, \{v_{1,t}, \dots, v_{i-1,t}\})$$

is indeed computable by the player when choosing the i^{th} vertex at round t , since the player has already observed $\mathcal{J}(\mathcal{A}_t, \{v_{1,t}, \dots, v_{i-1,t}\})$ after the first $i - 1$ source nodes are selected.

We then have the following result concerning the scaled pseudo-regret:

Theorem 11 (Symmetric loss, multiple sources). *Suppose $k > 1$ and the player uses the strategy $\mathcal{S}_{OSMD}^{sym,k}$ corresponding to the Online Greedy Algorithm with single-source strategy \mathcal{S}_{OSMD}^{sym} . Then the scaled pseudo-regret satisfies the bound*

$$\sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T^{(1-1/e)}(\mathcal{A}, \mathcal{S}_{OSMD}^{sym,k}) \leq 2^{\frac{1}{4}} k \sqrt{Tn}.$$

Comparing Theorem 11 to Theorem 10, we see an additional factor of k in the pseudo-regret upper bound. Similar results may be derived when alternative single-source strategies are used as subroutines in the Online Greedy Algorithm. We prove Theorem 11 and discuss properties of the general Online Greedy Algorithm in Appendix C.

4.1.3 Lower Bounds

We now establish lower bounds for the pseudo-regret in the case $k = 1$. This furnishes a better understanding of the hardness of the adversarial influence maximization problem. The general approach for deriving lower bounds is to produce a strategy for the adversary that forces the player to incur a certain level of regret regardless of which strategy is chosen.

The intrinsic difficulty of online influence maximization may vary widely depending on the topology of the underlying graph, and methods for deriving lower bounds may also differ accordingly. In the case of a complete graph, we have the following result:

Theorem 12. Suppose $G = \mathcal{K}_n$ is the complete graph on $n \geq 3$ vertices. Then the pseudo-regret satisfies the lower bound

$$\frac{2}{243}\sqrt{T} \leq \inf_{\mathcal{S} \in \mathcal{D}} \sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}).$$

Remark 3. Clearly, a gap exists between the lower bound derived in Theorem 12 and the upper bound appearing in Theorem 10. It is unclear which bound, if any, provides the proper minimax rate. However, note that if the lower bound were tight, it would imply that the proportion of vertices that the player misses by picking suboptimal source sets is constant, meaning the number of additional vertices the optimal source vertex infects is linear in the size of the graph. This differs substantially from the pseudo-regret of order \sqrt{n} known to be minimax optimal for the standard multi-armed bandit problem (and arises, for instance, in the case of directed graphs, as discussed in the next section).

The proof of Theorem 12, provided in Section D.1.1, studies the situation where the adversary selects the edges in a randomly chosen clique to be open at each time step.

4.2 Directed Graphs

We now derive upper and lower bounds for the pseudo-regret in the case of directed graphs, when $k = 1$.

4.2.1 Upper Bounds

The symmetric loss does not have a clear analog in the case of directed graphs. However, we may still use the node loss estimate for multi-armed bandit problems, given by equation (18). This leads to the following upper bound:

Theorem 13. Suppose the player uses the strategy $\mathcal{S}_{OSMD}^{node}$ corresponding to OSMD with the node loss $\hat{\ell}^{node}$ and appropriate parameters. Then the pseudo-regret satisfies the bound

$$\sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}_{OSMD}^{node}) \leq 2^{\frac{3}{2}} \sqrt{Tn}.$$

The proof follows from standard arguments (Bubeck and Cesa-Bianchi, 2012), so we do not state it here.

Remark 4. In the case $k > 1$, we may again use the Online Greedy Algorithm discussed in Section 4.1.2 to obtain a player strategy composed of parallel runs of a single-source strategy. If the player uses the single-source strategy $\mathcal{S}_{OSMD}^{node}$, we may obtain the scaled pseudo-regret bound

$$\sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T^{(1-1/e)}(\mathcal{A}, \mathcal{S}_{OSMD}^{node,k}) \leq 2^{\frac{3}{2}} k \sqrt{Tn}.$$

4.2.2 Lower Bounds

Finally, we provide a lower bound for the directed complete graph on n vertices. (This refers to the case where all edges are present and bidirectional.) We have the following result:

Theorem 14. *Suppose G is the directed complete graph on n vertices. Then the pseudo-regret satisfies the lower bound*

$$\frac{1}{48\sqrt{6}}\sqrt{nT} \leq \inf_{\mathcal{S} \in \mathcal{P}} \sup_{A \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}).$$

The proof of Theorem 14 is provided in Section D.1.2.

Notably, the lower bound in Theorem 14 matches the upper bound in Theorem 13, up to constant factors. Thus, the minimax pseudo-regret for the influence maximization problem is $\Theta(\sqrt{nT})$ in the case of directed graphs. In the case of undirected graphs, however (cf. Theorem 12), we were only able to obtain a pseudo-regret lower bound of $\Omega(\sqrt{T})$. This is due to the fact that in undirected graphs, one may learn about the loss of other nodes at time t besides the loss at \mathcal{S}_t . In contrast, it is possible to construct adversarial strategies for directed graphs that do not provide information regarding the loss incurred by choosing a source vertex other than \mathcal{S}_t .

Finally, we remark that a different choice of G might affect the lower bound, since influence maximization may be easier for some graph topologies than others. However, Theorem 14 shows that the case of the complete graph is always guaranteed to incur a pseudo-regret that matches the general upper bound in Theorem 13, implying that this is the minimax optimal rate for any class of graphs containing the complete graph.

5 Proofs

We now outline the proofs of our main results. We first consider the stochastic spreading model and prove our theorems for linear threshold and triggering models. Next, we prove the upper and lower bounds on pseudo-regret for adversarial models.

5.1 Proofs for Linear Threshold Models

We begin by proving the theorems in Section 3.2 concerning influence bounds for the linear threshold model.

5.1.1 Proof of Theorem 1

We first establish the following lemma, proved in Appendix A.1:

Lemma 1. *Define P_A to be set of paths from A to \bar{A} such that only the first vertex of the path belongs to A ; i.e.,*

$$P_A = A \cup \{p_1 p_2 \mid p_1 \in \mathcal{P}_{A \rightarrow i}^1, p_2 \in \mathcal{P}_{i \rightarrow \bar{A} | \bar{A}}, \text{ for some } i \in \bar{A}\}.$$

For any $A \subseteq V$, the influence of A in the linear threshold model is given by

$$\mathcal{I}(A) = \omega(P_A). \tag{20}$$

Applying Lemma 1, we have

$$\begin{aligned}
\mathcal{I}(A) &= \sum_{p \in P_A} \omega(p) \\
&= |A| + \sum_{i \in \bar{A}} \omega(\mathcal{P}_{A \rightarrow i}^1) \omega(\mathcal{P}_{i \rightarrow \bar{A}|\bar{A}}) \\
&= |A| + \sum_{i \in \bar{A}} \omega(\mathcal{P}_{A \rightarrow i}^1) \omega \left(\bigcup_{k \leq n-|A|-1} \mathcal{P}_{i \rightarrow \bar{A}|\bar{A}}^k \right) \\
&\leq |A| + \sum_{i \in \bar{A}} \omega(\mathcal{P}_{A \rightarrow i}^1) \omega \left(\bigcup_{k \leq n-|A|-1} \mathcal{W}_{i \rightarrow \bar{A}|\bar{A}}^k \right) \\
&= |A| + b_{\bar{A}}^T \sum_{k=0}^{n-|A|-1} B_{\bar{A},\bar{A}}^k \mathbf{1}_{\bar{A}} \\
&\leq |A| + b_{\bar{A}}^T \sum_{k=0}^{\infty} B_{\bar{A},\bar{A}}^k \mathbf{1}_{\bar{A}} \\
&\stackrel{(a)}{=} |A| + b_{\bar{A}}^T (I - B_{\bar{A},\bar{A}})^{-1} \mathbf{1}_{\bar{A}}.
\end{aligned} \tag{21}$$

Here, (a) is valid because $\rho(B_{\bar{A},\bar{A}}) < 1$, which holds since $B_{\bar{A},\bar{A}}$ is sub-stochastic (i.e., at least one column sum is strictly less than 1) as long as at least one edge exists from A to \bar{A} .

5.1.2 Proof of Theorem 2

The statement of the theorem is immediate from Lemma 1. We focus on establishing the special cases. For $k = 0$, we have $\omega(P_A^0) = |A|$. For $k \geq 1$, we have

$$\omega(P_A^k) = \sum_{i \in \bar{A}} \omega(\mathcal{P}_{A \rightarrow i}^1) \omega(\mathcal{P}_{i \rightarrow \bar{A}|\bar{A}}^{k-1}). \tag{22}$$

For $m = 1$, note that P_A^1 is simply the collection of edges from A to \bar{A} . The sum of the weights, which is $\omega(\mathcal{P}_{A \rightarrow \bar{A}}^1)$, equals $b_{\bar{A}}^T \mathbf{1}$. This yields inequality (8). For $m = 2$, equation (22) gives

$$\omega(P_A^2) = \sum_{i \in \bar{A}} \omega(\mathcal{P}_{A \rightarrow i}^1) \omega(\mathcal{P}_{i \rightarrow \bar{A}}^1) = \sum_{i \in \bar{A}} b_{\bar{A}}(i) B_{\bar{A},\bar{A}} \mathbf{1}_{\bar{A}}(i) = b_{\bar{A}}^T B_{\bar{A},\bar{A}} \mathbf{1}_{\bar{A}}.$$

Summing up over over $k = 1$ and $k = 2$ yields inequality (9). For $m = 3$, we use a similar idea and note that any path in P_A^3 must consist of an edge from A to a vertex $i \in \bar{A}$, followed by a path of length 2 in $\mathcal{P}_{i \rightarrow \bar{A}|\bar{A}}^2$. Thus,

$$\omega(\mathcal{P}_{i \rightarrow \bar{A}|\bar{A}}^2) = \omega(\mathcal{W}_{i \rightarrow \bar{A}|\bar{A}}) - \omega(\mathcal{W}_{i \rightarrow i|\bar{A}}^2) = B_{\bar{A}\bar{A}}^2 \mathbf{1}_{\bar{A}}(i) - B_{\bar{A},\bar{A}}^2(i, i).$$

Hence,

$$\begin{aligned}
\omega(P_A^3) &= \sum_{i \in \bar{A}} \omega(\mathcal{P}_{A \rightarrow i}^1) \omega(\mathcal{P}_{i \rightarrow \bar{A}|\bar{A}}^2) \\
&= \sum_{i \in \bar{A}} b_{\bar{A}}(i) \times \left(B_{\bar{A}\bar{A}}^2 \mathbf{1}_{\bar{A}}(i) - B_{\bar{A},\bar{A}}^2(i, i) \right) \\
&= b_{\bar{A}}^T (B_{\bar{A},\bar{A}}^2 - \text{Diag}(B_{\bar{A},\bar{A}}^2)) \mathbf{1}_{\bar{A}}.
\end{aligned}$$

Summing up the corresponding terms for $k = 1, 2$, and 3 yields inequality (10).

5.1.3 Proof of Theorem 3

We have

$$\begin{aligned}
UB - LB_1 &= b_{\bar{A}}^T \mathbf{1}_{\bar{A}} - b_{\bar{A}}^T (I - B_{\bar{A},\bar{A}})^{-1} \mathbf{1}_{\bar{A}} \\
&= b_{\bar{A}}^T \mathbf{1}_{\bar{A}} - b_{\bar{A}}^T (I + B_{\bar{A},\bar{A}} + B_{\bar{A},\bar{A}}^2 + \cdots) \mathbf{1}_{\bar{A}} \\
&= b_{\bar{A}}^T (B_{\bar{A},\bar{A}} + B_{\bar{A},\bar{A}}^2 + \cdots) \mathbf{1}_{\bar{A}} \\
&\leq \|b_{\bar{A}}\|_1 \left\| (B_{\bar{A},\bar{A}} + B_{\bar{A},\bar{A}}^2 + \cdots) \mathbf{1}_{\bar{A}} \right\|_{\infty} \\
&\leq LB_1 \left\| (B_{\bar{A},\bar{A}} + B_{\bar{A},\bar{A}}^2 + \cdots) \right\|_{\infty, \infty} \|\mathbf{1}_{\bar{A}}\|_{\infty} \\
&\leq LB_1 \times \frac{\lambda_{\bar{A}, \infty}}{1 - \lambda_{\bar{A}, \infty}},
\end{aligned}$$

so

$$\frac{UB}{LB_1} \leq \frac{1}{1 - \lambda_{\bar{A}, \infty}}.$$

For LB_2 , we have

$$\begin{aligned}
UB - LB_2 &= b_{\bar{A}}^T (B_{\bar{A},\bar{A}}^2 + B_{\bar{A},\bar{A}}^3 + \cdots) \mathbf{1}_{\bar{A}} \\
&= b_{\bar{A}}^T (I + B_{\bar{A},\bar{A}}) (B_{\bar{A},\bar{A}}^2 + B_{\bar{A},\bar{A}}^4 + \cdots) \mathbf{1}_{\bar{A}} \\
&\leq \|b_{\bar{A}}^T (I + B_{\bar{A},\bar{A}})\|_1 \left\| (B_{\bar{A},\bar{A}}^2 + B_{\bar{A},\bar{A}}^4 + \cdots) \mathbf{1}_{\bar{A}} \right\|_{\infty} \\
&\leq LB_2 \left\| (B_{\bar{A},\bar{A}}^2 + B_{\bar{A},\bar{A}}^4 + \cdots) \right\|_{\infty, \infty} \|\mathbf{1}_{\bar{A}}\|_{\infty} \\
&\leq LB_2 \cdot \left(\frac{\lambda_{\bar{A}, \infty}^2}{1 - \lambda_{\bar{A}, \infty}^2} \right),
\end{aligned}$$

so

$$\frac{UB}{LB_2} \leq \frac{1}{1 - \lambda_{\bar{A}, \infty}^2}.$$

5.2 Proofs for Triggering Models

In this section, we provide the proofs for the theorems in Section 3.3 involving general triggering models.

5.2.1 Proof of Theorem 4

Let X_i be the indicator random variable for the event “vertex i is infected,” and for any path p , let Y_p be the indicator random variable for the event “all edges in p are live.” Recall that $P_{A \rightarrow i}$ is the set of all paths from A to i , such that only the starting vertex belongs to A . We now use the representation

$$X_i = 1 - \prod_{p \in P_{A \rightarrow i}} (1 - Y_p).$$

(Such a representation is considered briefly in the proof of Lemma 6 in [Lemonnier et al., 2014](#), but it is not used for anything beyond the correlation proof.) Since

$$1 - \sum_{p \in P_{A \rightarrow i}} Y_i \leq \prod_{p \in P_{A \rightarrow i}} (1 - Y_p),$$

we have the upper bound

$$\mathbb{E}X_i \leq \sum_{p \in P_{A \rightarrow i}} \mathbb{E}Y_p = \sum_{p \in P_{A \rightarrow i}} \omega(p).$$

Note that this bound is sharp when there is at most one path in $P_{A \rightarrow i}$. Summing up over all vertices, we have

$$\mathcal{I}(A) \leq \sum_i \sum_{p \in P_{A \rightarrow i}} \omega(p) = \sum_{p \in P_A} \omega(p).$$

Proceeding exactly as in the derivation of inequality (21) in the proof of Theorem 1, we then obtain the desired bound.

5.2.2 Proof of Theorem 5

Let X_i be the indicator random variable for the event “vertex i is infected,” and for any path p , let Y_p be the indicator random variable for the event “all edges in p are live.” As in the proof of Theorem 4, we use the representation (39). Taking expectations in equation (39), we have

$$\begin{aligned} \mathbb{E}[X_i] &= 1 - \mathbb{E} \left[\prod_{p \in P_{A \rightarrow i}} (1 - Y_p) \right] \\ &\stackrel{(a)}{\geq} 1 - \inf_{p \in P_{A \rightarrow i}} \mathbb{E}(1 - Y_p) \\ &= \sup_{p \in P_{A \rightarrow i}} \mathbb{E}Y_p \\ &= \sup_{p \in P_{A \rightarrow i}} \omega(p), \end{aligned}$$

where (a) follows from the fact that

$$1 - Y_q \geq \prod_{p \in P_{A \rightarrow i}} (1 - Y_p)$$

for any q in $P_{A \rightarrow i}$. This completes the proof.

5.2.3 Proof of Theorem 6

Using simple algebraic manipulations, we have

$$\begin{aligned}
\mathcal{I}(A) &\leq |A| + b_{\bar{A}}^T \sum_{k=0}^{n-|A|-1} B_{\bar{A},\bar{A}}^k \mathbf{1}_{\bar{A}} \\
&\leq |A| + \|b_{\bar{A}}\|_1 \left\| \left(\sum_{k=0}^{n-|A|-1} B_{\bar{A},\bar{A}}^k \right) \mathbf{1}_{\bar{A}} \right\|_{\infty} \\
&\stackrel{(a)}{\leq} |A| + \lambda_{\infty} |A| \left\| \left(\sum_{k=0}^{n-|A|-1} B_{\bar{A},\bar{A}}^k \right) \right\|_{\infty, \infty} \|\mathbf{1}_{\bar{A}}\|_{\infty} \\
&\stackrel{(b)}{\leq} |A| + \lambda_{\infty} |A| \left\| \left(\sum_{k=0}^{n-|A|-1} B^k \right) \right\|_{\infty, \infty} \\
&\leq |A| + \lambda_{\infty} |A| \cdot \frac{1 - \lambda_{\infty}^{n-|A|}}{1 - \lambda_{\infty}}. \tag{23}
\end{aligned}$$

In step (a) in the above derivation, we take $\lambda_{\infty} = \|B\|_{\infty, \infty}$ and deduce that the total weight of outgoing edges from A to \bar{A} , which equals $\|b_{\bar{A}}\|_1$, is at most $\lambda_{\infty} |A|$. In step (b), we use the fact that B is a positive matrix of which $B_{\bar{A},\bar{A}}$ is a submatrix, so $\|B_{\bar{A},\bar{A}}^k\|_{\infty, \infty} \leq \|B^k\|_{\infty, \infty}$, for all $k \geq 1$.

If $\lambda_{\infty} < 1$, we may simplify inequality (23) to obtain

$$\mathcal{I}(A) \leq |A| + |A| \cdot \frac{\lambda_{\infty}}{1 - \lambda_{\infty}} = \frac{|A|}{1 - \lambda_{\infty}}.$$

5.3 Upper Bounds for Adversarial Models

In this section, we prove our upper bounds for online influence maximization. To this end, we analyze the Exp3 algorithm and then describe the OSMD algorithm.

5.3.1 Exp3 Algorithms

Theorem 9 follows from a more general proposition:

Proposition 3. *Suppose the loss estimates $\{\hat{\ell}_{i,t}\}$ satisfy*

$$\mathbb{E}_{\mathcal{S}_t \sim p_t} \hat{\ell}_{i,t} = \ell_{i,t}, \tag{24}$$

$$\mathbb{E}_{I \sim p_t} \hat{\ell}_{I,t} = \ell_{\mathcal{S}_t,t}, \tag{25}$$

$$\mathbb{E}_{\mathcal{S}_t \sim p_t} \mathbb{E}_{I \sim p_t} \hat{\ell}_{I,t}^2 \leq v(n), \tag{26}$$

for some function $v(n)$, where the expectation in equation (25) is taken conditioned on \mathcal{S}_t . Then the strategy \mathcal{S} corresponding to the Exp3 algorithm with loss $\hat{\ell}$ and $\eta_t = \sqrt{\frac{2 \log n}{T v(n)}}$ satisfies the pseudo-regret bound

$$\sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}) \leq \sqrt{2T v(n) \log n}.$$

Proof. We follow the proof technique of Theorem 3.1 in [Bubeck and Cesa-Bianchi \(2012\)](#). Our goal is to prove that for the player strategy \mathcal{S} governed by the Exp3 algorithm, and for any adversarial strategy $\mathcal{A} \in \mathcal{A}$, the following statement holds:

$$\bar{R}_T(\mathcal{A}, \mathcal{S}) \leq \frac{v(n)\eta T}{2} + \frac{\log n}{\eta}, \quad (27)$$

where $\eta = \sqrt{\frac{2 \log n}{Tv(n)}}$. The desired bound then follows from simple algebra.

We begin by writing

$$\sum_{t=1}^n \ell_{\mathcal{S}_t, t} - \sum_{t=1}^n \ell_{k, t} = \sum_{t=1}^n \mathbb{E}_{I \sim p_t} \hat{\ell}_{I, t} - \sum_{t=1}^n \ell_{k, t}. \quad (28)$$

Note that taking an expectation with respect to \mathcal{A} and \mathcal{S} , followed by a maximum over $1 \leq k \leq n$, yields the pseudo-regret $\bar{R}_T(\mathcal{A}, \mathcal{S})$.

We expand the first term in equation (28):

$$\mathbb{E}_{I \sim p_t} \hat{\ell}_{I, t} = \frac{1}{\eta} \log \mathbb{E}_{I \sim p_t} \exp \left(-\eta (\hat{\ell}_{I, t} - \mathbb{E}_{I' \sim p_t} \hat{\ell}_{I', t}) \right) - \frac{1}{\eta} \log \mathbb{E}_{I \sim p_t} \exp(-\eta \hat{\ell}_{I, t}) := A + B, \quad (29)$$

where I' is an independent copy of I . Using the inequalities $\log x \leq x - 1$ and

$$\exp(-x) - 1 + x \leq \frac{x^2}{2},$$

we then have

$$\begin{aligned} A &= \frac{1}{\eta} \log \mathbb{E}_{I \sim p_t} \exp(-\eta \hat{\ell}_{I, t}) + \frac{1}{\eta_t} \mathbb{E}_{I' \sim p_t} [\eta \hat{\ell}_{I', t}] \\ &\leq \frac{1}{\eta} \mathbb{E}_{I \sim p_t} \left[\exp(-\eta \hat{\ell}_{I, t}) - 1 + \eta \hat{\ell}_{I, t} \right] \\ &\leq \frac{1}{\eta} \mathbb{E}_{I \sim p_t} \frac{\eta^2 \hat{\ell}_{I, t}^2}{2} \\ &= \frac{\eta}{2} \mathbb{E}_{I \sim p_t} \hat{\ell}_{I, t}^2. \end{aligned}$$

Furthermore, recalling that $\hat{L}_{i, t} = \hat{L}_{i, t-1} + \hat{\ell}_{i, t}$ and $p_{i, t} = \frac{\exp(-\eta \hat{L}_{i, t-1})}{\sum_{k=1}^n \exp(-\eta \hat{L}_{k, t-1})}$, we have

$$B = -\frac{1}{\eta} \log \frac{\sum_{i=1}^n \exp(-\eta \hat{L}_{i, t})}{\sum_{i=1}^n \exp(-\eta \hat{L}_{i, t-1})} = \Phi_{t-1}(\eta) - \Phi_t(\eta),$$

where we have used the shorthand notation

$$\Phi_t(\eta) := \frac{1}{\eta} \log \frac{1}{n} \sum_{i=1}^n \exp(-\eta \hat{L}_{i, t}).$$

Substituting the bounds back into equations (28) and (29), we then obtain

$$\begin{aligned} \sum_{t=1}^T \ell_{\mathcal{S}_t, t} - \sum_{t=1}^T \ell_{k, t} &\leq \frac{\eta}{2} \sum_{t=1}^T \mathbb{E}_{I \sim p_t} \hat{\ell}_{I, t}^2 + \sum_{t=1}^T (\Phi_{t-1}(\eta) - \Phi_t(\eta)) - \sum_{t=1}^T \ell_{k, t} \\ &= \frac{\eta}{2} \sum_{t=1}^T \mathbb{E}_{I \sim p_t} \hat{\ell}_{I, t}^2 - \Phi_T(\eta) - \sum_{t=1}^T \ell_{k, t}, \end{aligned} \quad (30)$$

where the last equality uses the fact that $\Phi_0 \equiv 0$. Furthermore,

$$\begin{aligned} -\Phi_T(\eta) &= \frac{\log n}{\eta} - \frac{1}{\eta} \log \left(\sum_{i=1}^n \exp(-\eta \widehat{L}_{i,T}) \right) \\ &\leq \frac{\log n}{\eta} - \frac{1}{\eta} \log(\exp(-\eta \widehat{L}_{k,T})) \\ &= \frac{\log n}{\eta} + \sum_{t=1}^T \widehat{\ell}_{k,t}. \end{aligned}$$

Substituting back into inequality (30) and taking an expectation with respect to \mathcal{S}_t , we conclude that

$$\mathbb{E}_{\mathcal{S}_t \sim p_t} \sum_{t=1}^T \ell_{\mathcal{S}_t,t} - \sum_{t=1}^T \ell_{k,t} \leq \mathbb{E} \left[\frac{\eta}{2} \sum_{t=1}^T \mathbb{E}_{I \sim p_t} \widehat{\ell}_{I,t}^2 \right] + \frac{\log n}{\eta} \leq \frac{v(n)\eta T}{2} + \frac{\log n}{\eta}.$$

Since this inequality holds for all $1 \leq k \leq n$, we arrive at the desired bound (27). \square

As a result of Proposition 3, we only need to verify the statements (24), (25), and (26), for $v(n) = \frac{n+1}{2}$, in the setting of Theorem 9. Details are provided in Appendix B.1.

5.3.2 OSMD Algorithms

We now outline the OSMD algorithm, which generates a sequence of probability distributions $\{p_t\}$ to be employed by the player on successive rounds. Let $\Delta^n \subseteq \mathbb{R}^n$ denote the probability simplex.

Online Stochastic Mirror Descent (OSMD) with loss estimates $\{\widehat{\ell}_{i,t}\}$

Given: A Legendre function F defined on \mathbb{R}^n , with associated Bregman divergence

$$D_F(p, q) = F(p) - F(q) - (p - q)^T \nabla F(q),$$

and a learning rate $\eta > 0$.

Output: A stochastic player strategy $\{\mathcal{S}_t\}$.

Let $p_1 \in \arg \min_{p \in \Delta^n} F(p)$.

For each round $t = 1, \dots, T$:

- (1) Draw a vertex \mathcal{S}_t from the distribution p_t .
- (2) Compute the vector of loss estimates $\widehat{\ell}_t = \{\widehat{\ell}_{i,t}\}$.
- (3) Set $w_{t+1} = \nabla F^* \left(\nabla F(p_t) - \eta \widehat{\ell}_t \right)$, where F^* is the convex conjugate of F .
- (4) Compute the new distribution $p_{t+1} = \arg \min_{p \in \Delta^n} D_F(p, w_{t+1})$.

In general, the OSMD algorithm is defined with respect to a compact, convex set $\mathcal{K} \subseteq \mathbb{R}^n$. The updates are characterized by noisy estimates of the gradient of the loss function, which we may

conveniently define to be $\widehat{\ell}_t$ in the present scenario. For more details and generalizations, we refer the reader to [Bubeck and Cesa-Bianchi \(2012\)](#).

We will use the following result:

Proposition 4 (Theorem 5.10 of [Bubeck and Cesa-Bianchi, 2012](#)). *Suppose the loss functions $\{\ell_{i,t}\}$ are nonnegative and bounded by 1. The strategy \mathcal{S} corresponding to the OSMD algorithm with loss estimates $\widehat{\ell}$, learning rate $\eta > 0$, and Legendre function F_ψ , where ψ is a 0-potential, satisfies the pseudo-regret bound*

$$\sup_{\mathcal{A} \in \mathcal{A}} \overline{R}_T(\mathcal{A}, \mathcal{S}) \leq \frac{\sup_{p \in \Delta^n} F_\psi(p) - F_\psi(p_1)}{\eta} + \frac{\eta}{2} \sum_{t=1}^T \sum_{i=1}^n \mathbb{E} \left[\frac{\widehat{\ell}_{i,t}^2}{(\psi^{-1})'(p_{i,t})} \right].$$

We will formally define 0-potentials and the associated Legendre functions in [Appendix B.2](#). In the analysis in this paper, we will take $\psi(x) = \frac{1}{x^2}$, yielding the Legendre function $F_\psi(x) = -2 \sum_{i=1}^n x_i^{1/2}$. The pseudo-regret bound in [Proposition 4](#) may then be analyzed and bounded accordingly in various settings of interest. Details for the proof of [Theorem 10](#) are provided in [Appendix B.2.2](#).

5.4 Lower Bounds for Adversarial Models

We now turn to establishing lower bounds for adversarial influence maximization. The proofs of [Theorems 12](#) and [14](#) are based on the same general strategy, which is summarized in the following proposition. To unify our results with standard bandit notation ([Bubeck and Cesa-Bianchi, 2012](#)), we use the shorthand

$$X_{i,t} = f(\mathcal{A}_t, \{i\})$$

to denote the reward incurred at time t when the player chooses $\mathcal{S}_t = \{i\}$. Then

$$\overline{R}_T(\mathcal{A}, \mathcal{S}) = \max_{1 \leq i \leq n} \mathbb{E}_{\mathcal{A}, \mathcal{S}} \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}).$$

Proposition 5. *Consider a deterministic player strategy $\mathcal{S} \in \mathcal{P}_d$. Let $\mathcal{A}^0, \mathcal{A}^1, \dots, \mathcal{A}^n$ be stochastic adversarial strategies such that for each \mathcal{A}^i , the set of edges played at time t is independent of the past actions of the adversary. Let $\mathbb{P}_0, \mathbb{P}_1, \dots, \mathbb{P}_n$ denote the corresponding measures on the feedback \mathcal{J}^T , allowing for possible randomization only in the strategy of the adversary. Let \mathbb{E}_i denote the expectation with respect to \mathbb{P}_i . Suppose*

$$r \leq \min_{j \neq i} \mathbb{E}_i [X_{i,t} - X_{j,t}], \quad \forall 1 \leq t \leq T, \quad (31)$$

and

$$\sum_{i=1}^n KL(\mathbb{P}_0, \mathbb{P}_i) \leq D. \quad (32)$$

Then

$$rT \left(\frac{n-1}{n} - \sqrt{\frac{D}{2n}} \right) \leq \frac{1}{n} \sum_{i=1}^n \mathbb{E}_i \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}). \quad (33)$$

In particular, if the bounds [\(31\)](#) and [\(32\)](#) hold uniformly for all choices of $\mathcal{S} \in \mathcal{P}_d$, then

$$rT \left(\frac{n-1}{n} - \sqrt{\frac{D}{2n}} \right) \leq \inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}} \overline{R}_T(\mathcal{A}, \mathcal{S}). \quad (34)$$

Remark 5. We remark briefly about the roles of the strategies $\{\mathcal{A}^i\}$ appearing in Proposition 5. In practice, the strategies are chosen to be similar, except selecting i as the source node is slightly more advantageous when the adversary uses strategy \mathcal{A}^i . The strategy \mathcal{A}^0 is a baseline strategy that treats all nodes identically. Thus, the lower bound provided by Proposition 5 is the product of the cost of an incorrect choice of the source vertex, given by r , and a factor that determines how easy it is to distinguish the adversary strategies from each other, which depends on D .

Proof. We follow the method used in the proof of Theorem 3.5 in [Bubeck and Cesa-Bianchi \(2012\)](#). We first show how to obtain the bound (34) from the set of uniform bounds (33). Note that for any $\mathcal{S} \in \mathcal{P}$, we have

$$\begin{aligned}
\sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}) &= \sup_{\mathcal{A} \in \mathcal{A}} \max_{1 \leq i \leq n} \mathbb{E}_{\mathcal{A}, \mathcal{S}} \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}) \\
&\geq \max_{1 \leq j \leq n} \max_{1 \leq i \leq n} \mathbb{E}_{\mathcal{S}} \mathbb{E}_{\mathcal{A}^j} \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}) \\
&= \max_{1 \leq j \leq n} \max_{1 \leq i \leq n} \mathbb{E}_{\mathcal{S}} \mathbb{E}_j \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}) \\
&\geq \max_{1 \leq i \leq n} \mathbb{E}_{\mathcal{S}} \mathbb{E}_i \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}) \\
&\geq \mathbb{E}_{\mathcal{S}} \left[\frac{1}{n} \sum_{i=1}^n \mathbb{E}_i \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}) \right],
\end{aligned}$$

where we have used the fact that the maximum is at least as large as the average in the final inequality. Since any player strategy in \mathcal{P} lies in the convex hull of deterministic player strategies, a uniform bound (33) over \mathcal{P}_d implies that inequality (34) holds, as well.

We now turn to the proof of inequality (33). The idea is to show that on average, the player incurs a certain loss whenever the wrong source vertex is played, and this event must happen sufficiently often. We first write

$$\begin{aligned}
\mathbb{E}_i \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}) &= \sum_{t=1}^T \mathbb{E}_i \left[\sum_{j \neq i} (X_{i,t} - X_{j,t}) \mathbf{1}_{\{\mathcal{S}_t = \{j\}\}} \right] \\
&= \sum_{j \neq i} \sum_{t=1}^T \mathbb{E}_i [X_{i,t} - X_{j,t}] \mathbb{E}_i [\mathbf{1}_{\{\mathcal{S}_t = \{j\}\}}].
\end{aligned}$$

In the last equality, we have used the assumption that the adversary's action at each time is independent of the past to conclude that the difference in rewards $X_{i,t} - X_{j,t}$ (which depends on the adversary's action at time t) is independent of the indicator $\mathbf{1}_{\{\mathcal{S}_t = \{j\}\}}$ (which depends on the sequence of feedback received up to time $t - 1$). Using the bound (31), it follows that

$$\mathbb{E}_i \sum_{t=1}^T (X_{i,t} - X_{\mathcal{S}_t,t}) = \sum_{j \neq i} r \mathbb{E}_i [T_j],$$

where $T_i = |\{t : \mathcal{S}_t = \{i\}\}|$ denotes the number of times vertex i is selected as the source.

Now let U_T denote a vertex drawn according to the distribution $q_T = (q_{1,T}, \dots, q_{n,T})$, where $q_{i,T} = \frac{T_i}{T}$. The derivations above imply that

$$\mathbb{E}_i \sum_{t=1}^T (X_{i,t} - X_{S_t,t}) = rT \sum_{j \neq i} \mathbb{P}_i \{U_T = j\} = rT (1 - \mathbb{P}_i \{U_T = i\}),$$

so

$$\frac{1}{n} \sum_{i=1}^n \mathbb{E}_i \sum_{t=1}^T (X_{i,t} - X_{S_t,t}) = rT \left(1 - \frac{1}{n} \sum_{i=1}^n \mathbb{P}_i \{U_T = i\} \right). \quad (35)$$

By Pinsker's inequality, we have

$$\mathbb{P}_i \{U_T = i\} \leq \mathbb{P}_0 \{U_T = i\} + \sqrt{\frac{1}{2} KL(\mathbb{P}'_0, \mathbb{P}'_i)},$$

where \mathbb{P}'_i denotes the distribution of U_T under the adversarial strategy \mathcal{A}^i . By Jensen's inequality, we therefore have

$$\frac{1}{n} \sum_{i=1}^n \mathbb{P}_i \{U_T = i\} \leq \frac{1}{n} + \frac{1}{n} \sum_{i=1}^n \sqrt{\frac{1}{2} KL(\mathbb{P}'_0, \mathbb{P}'_i)} \leq \frac{1}{n} + \sqrt{\frac{1}{2n} \sum_{i=1}^n KL(\mathbb{P}'_0, \mathbb{P}'_i)}. \quad (36)$$

Finally, the chain rule for KL divergence implies that

$$KL(\mathbb{P}'_0, \mathbb{P}'_i) = KL(\mathbb{P}_0, \mathbb{P}_i) + \sum_{\mathcal{J}^T} \mathbb{P}_0 \{\mathcal{J}^T\} KL(\mathbb{P}'_0 \{\cdot | \mathcal{J}^T\}, \mathbb{P}'_i \{\cdot | \mathcal{J}^T\}). \quad (37)$$

Note that conditional on \mathcal{J}^T , the distribution of U_T is the same under \mathbb{P}'_0 and \mathbb{P}'_i , since the player uses a deterministic strategy. Thus, equation (37) implies that

$$\sum_{i=1}^n KL(\mathbb{P}'_0, \mathbb{P}'_i) = \sum_{i=1}^n KL(\mathbb{P}_0, \mathbb{P}_i) \leq D. \quad (38)$$

Combining inequalities (35), (36), and (38), we arrive at the desired result (33). \square

To prove Theorems 12 and 14, it thus remains to find an appropriate set of strategies $\{\mathcal{A}^0, \mathcal{A}^1, \dots, \mathcal{A}^n\}$ and verify the bounds (31) and (32). Details for the proofs are provided in Appendix D.1.

6 Simulations

In this section, we report the results of various simulations in the stochastic spreading setting. For the first set of simulations, we generated an Erdős-Renyi graph with 900 vertices and edge probability $\frac{2}{n}$; a preferential attachment graph with 900 vertices, 10 initial vertices, and 3 edges for each added vertex; and a 30×30 grid. We generated 33 instances of edge probabilities for each graph, as follows: For each instance and each vertex i , we chose $\gamma(i)$ uniformly in $[\gamma_{\min}, 0.8]$, where γ_{\min} ranged from 0.0075 to 0.75 in increments of 0.0075. The probability that the incoming edge was chosen was $\frac{1-\gamma}{d(i)}$, where $d(i)$ is the degree of i . An initial infection set A of size 10 was chosen at random, and 50 simulations of the infection process were run to estimate the true influence. The upper and lower bounds and value of $\mathcal{I}(A)$ computed via simulations are shown in Figure 1. Note that the gap

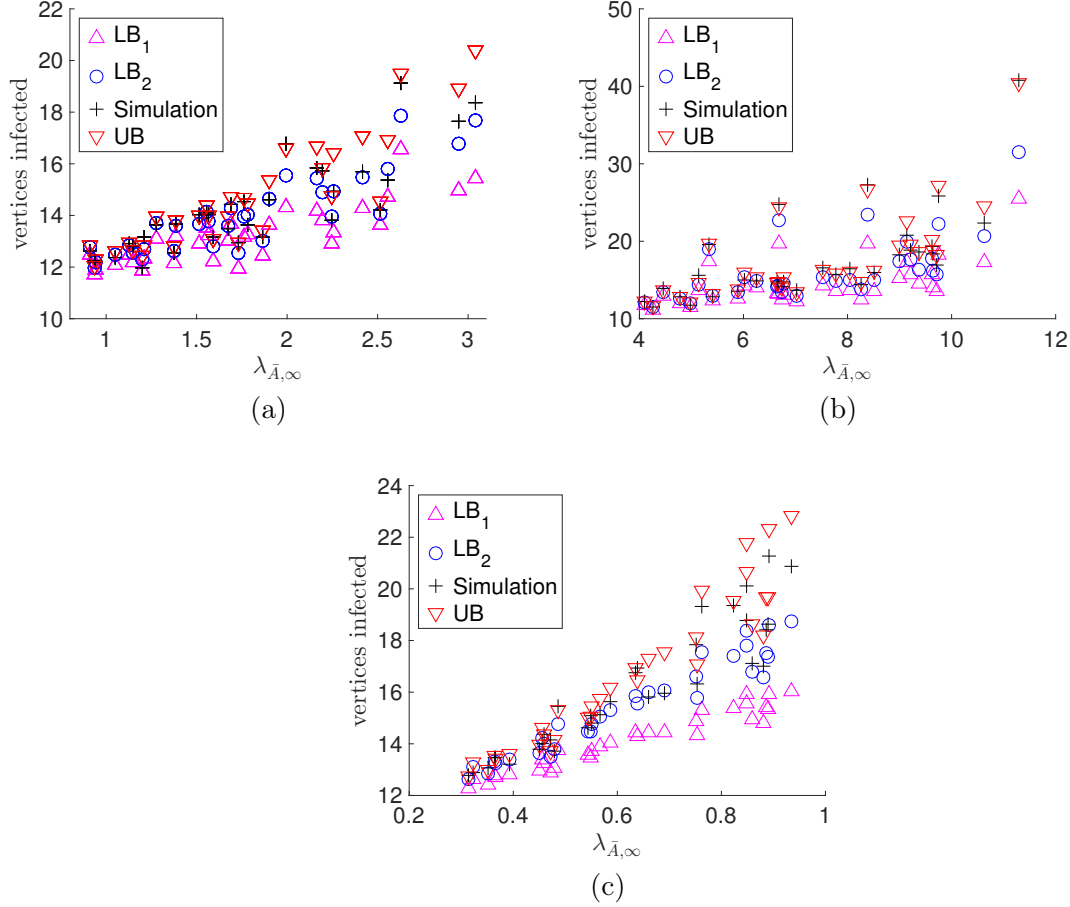


Figure 1: Lower bounds, upper bounds, and simulated influence for (a) Erdős-Rényi, (b) preferential attachment, and (c) 2D-grid graphs. For small values of $\lambda_{\bar{A},\infty}$, our bounds are tight.

between the upper and lower bounds indeed controlled for smaller values of $\lambda_{\bar{A},\infty}$, agreeing with the predictions of Theorem 3.

For the second set of simulations, we generated 10 of each of the following graphs: an Erdős-Rényi graph with 100 vertices and edge probability $\frac{2}{n}$; a preferential attachment graph with 100 vertices, 10 initial vertices, and 3 additional edges for each added vertex; and a grid graph with 100 vertices. For each of the 10 realizations, we also picked a value of $\gamma(i)$ for each vertex i uniformly in $[0.075, 0.8]$. The corresponding edge probabilities were assigned as before. We then selected sets A of size 10 using greedy algorithms to maximize LB_1 , LB_2 , and UB , as well as the estimated influence based on 50 simulated infections. Finally, we used 200 simulations to approximate the actual influence of each resulting set. The average influences, along with the average influence of a uniformly random subset of vertices of size 10, are plotted in Figure 2. Note that the greedy algorithms all perform comparably, although the sets selected using LB_2 and UB appear slightly better. The fact that the algorithm that uses UB performs well is somewhat unsurprising, since it takes into account the influence from all paths. However, note that maximizing UB does not lead to the theoretical guarantees we have derived for LB_1 and LB_2 . In Table 1, we report the runtimes scaled by the runtime of the LB_1 algorithm. As expected, the LB_1 algorithm is fastest, and the other algorithms may be much slower.

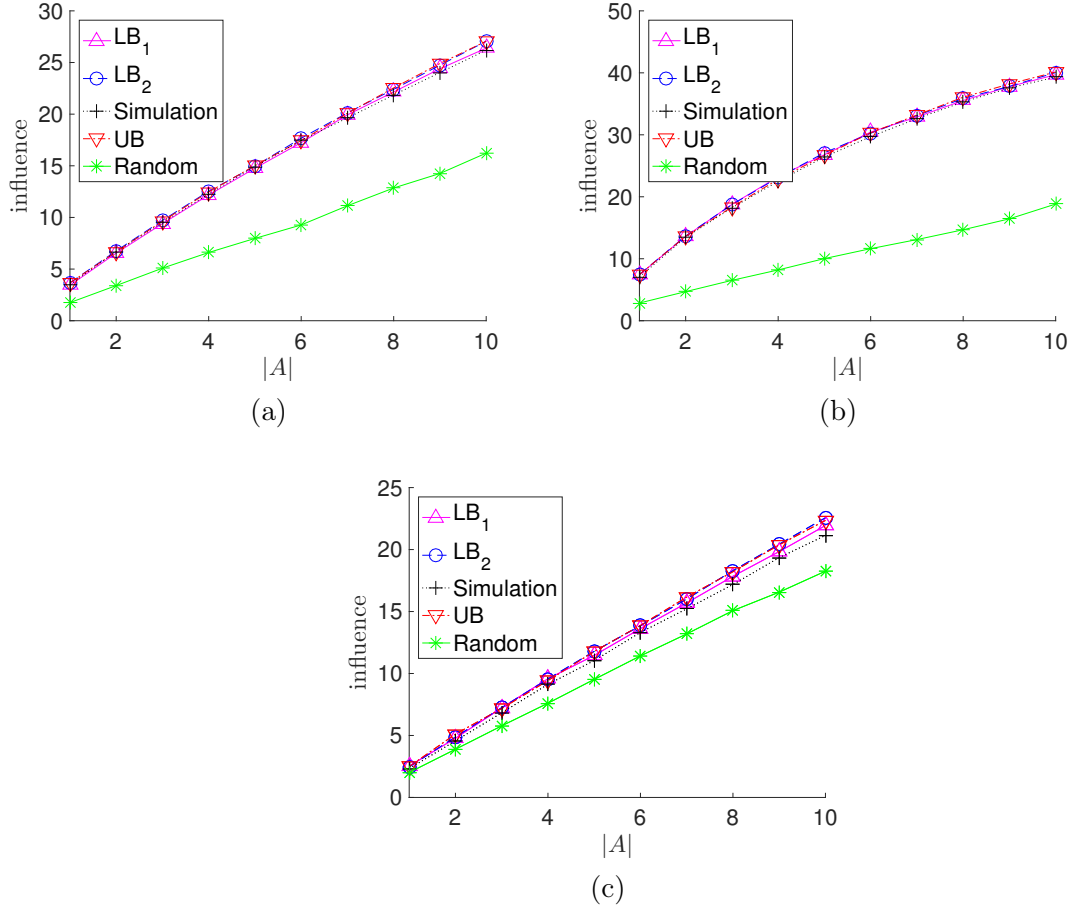


Figure 2: Simulated influence for sets $|A|$ selected by greedy algorithms and uniformly at random on (a) Erdős-Renyi, (b) preferential attachment, and (c) $2D$ -grid graphs. All greedy algorithms perform similarly, but the algorithms maximizing the simulated influence and UB are much more computationally intensive.

	LB_1	LB_2	UB	Simulation
Erdős-Renyi	1.00	2.36	27.43	710.58
Preferential attachment	1.00	2.56	28.49	759.83
$2D$ -grid	1.00	2.43	47.08	1301.73

Table 1: Runtimes for the influence maximization algorithms, scaled by the runtime of the greedy LB_1 algorithm. The corresponding lower bounds are much easier to compute, allowing for faster algorithms.

7 Discussion

We have proposed and analyzed algorithms for influence maximization in stochastic and adversarial contagion settings. Our results for stochastic spreading models are based on novel upper and lower bounds on the influence function that hold for general triggering models. Furthermore, the submodularity of our lower bounds guarantees that a sequential greedy maximization algorithm

produces a source set performing within a constant factor of the optimum. The resulting influence maximization algorithm is highly scalable and possesses attractive theoretical properties. In the case of adversarial spreading, we have devised player strategies that control the pseudo-regret uniformly across all possible oblivious adversarial strategies. For the problem of single-source influence maximization in complete networks, we have also derived minimax lower bounds that establish the fundamental hardness of the online influence maximization problem. In particular, our lower and upper bounds match up to constant factors in the case of undirected complete graphs, implying that our proposed player strategy is in some sense optimal. Although the stochastic and adversarial spreading models are appreciably different, our common goal in both settings is to obtain computationally efficient methods for selecting source sets of a certain cardinality with provable approximation guarantees in relation to the performance of the (uncomputable) optimal set.

Our work inspires a number of interesting questions for future study, several of which we now describe. An important open question concerns the closeness of our upper and lower bounds for the influence function in the case of general triggering models. Due to the complicated nature of the general lower bound LB_{trig} , it is not straightforward to characterize the proximity of LB_{trig} to the corresponding upper bound as a function of the edge weights and graph topology. Secondly, our results for stochastic spreading models may be extended via the conditional expectation decomposition employed by [Lee et al. \(2016\)](#) to generate sharper influence bounds for certain graph topologies, and it would be interesting to derive theoretical guarantees for the quality of improvement in such cases. Another worthwhile direction would be to derive theoretical guarantees for non-greedy algorithms in the lower bound maximization problem. As noted in [Angell and Schoenebeck \(2016\)](#), empirical studies show that the influence function of many real-world epidemics are *not* exactly submodular, necessitating appropriate modifications of the stochastic triggering models and occluding the theory for sequential greedy procedures. In the adversarial spreading setting, important open questions include closing the gap between upper and lower bounds on the minimax pseudo-regret in the case of undirected graphs; establishing lower bounds for other network topologies and multiple-source settings; and studying the case of non-oblivious adversaries. Our results only address a small subset of problems that may be posed and answered concerning a bandit theory of adversarial influence maximization with edge-level feedback.

A Additional Proofs for Stochastic Influence Maximization

In this Appendix, we prove Lemma [1](#). Additionally, we prove the submodularity of $\{LB_m\}_{m \geq 1}$ and LB_{trig} .

A.1 Proof of Lemma [1](#)

We use the reachability via live-edge paths interpretation of the linear threshold model. For a vertex $i \in \bar{A}$, let $P_{A \rightarrow i}$ denote the paths in P_A ending at i . Let X_i be the indicator random variable for the event “vertex i is infected.” For any path p , let Y_p be the indicator random variable for the event “all edges in p are active” (for a path of length 0, we take $Y_p = 1$). For any i , we have the equality

$$X_i = 1 - \prod_{p \in P_{A \rightarrow i}} (1 - Y_p) \stackrel{(a)}{=} 1 - \left(1 - \sum_{p \in P_{A \rightarrow i}} Y_p \right) = \sum_{p \in P_{A \rightarrow i}} Y_p, \quad (39)$$

where (a) holds because at most one live path may exist that reaches vertex i . Taking expectations on both sides and summing up, we have

$$\sum_{i \in \bar{A}} \mathbb{E}X_i = \sum_{i \in \bar{A}} \sum_{p \in P_{A \rightarrow i}} \mathbb{E}Y_p = \sum_{p \in P_A} \mathbb{E}Y_p = \sum_{p \in P_A} \omega(p).$$

A.2 Monotonicity and Submodularity

In this Appendix, we provide proofs of the theorems in Section 3.4 concerning monotonicity and submodularity of bounds for influence maximization.

A.2.1 Proof of Theorem 7

We first define some notation. For disjoint sets $A, B \subseteq V$, let $A + B$ denote the union $A \cup B$. If $B = \{x\}$, we also write $A + x$ for $A + \{x\}$. For sets $A_0, A_1, \dots, A_r \in V$, we denote the set of all paths $\{v_0, v_1, \dots, v_r\}$ of length r , such that $v_i \in A_i$ for $0 \leq i \leq r$, by

$$A_0 \rightarrow A_1 \rightarrow A_2 \rightarrow \dots \rightarrow A_r.$$

The combined weight of these paths is denoted by

$$\omega(A_0 \rightarrow A_1 \rightarrow A_2 \rightarrow \dots \rightarrow A_r).$$

The set of all possible paths from A to B , such that only the starting vertex lies in A , is denoted by $A \Rightarrow B$, and the weight of these paths is denoted by $\omega(A \Rightarrow B)$.

We begin by establishing monotonicity:

Lemma 2. *For each $m \geq 1$, the function LB_m is monotonically increasing.*

Proof. We will show that for any $A \subseteq V$ and $x \notin A$, the following inequality holds:

$$LB_m(A + x) \geq LB_m(A).$$

Let $B = V \setminus (A + x)$. We may express $LB_m(A)$ as follows:

$$LB_m(A) = |A| + \sum_{i=1}^m \omega(A \rightarrow B^i) + \sum_{\substack{i,j \geq 0 \\ i+j \leq m-1}} \omega(A \rightarrow B^i \rightarrow x \rightarrow B^j), \quad (40)$$

where we use the shorthand

$$B^i := \underbrace{B \rightarrow B \rightarrow \dots \rightarrow B}_{i \text{ times.}}$$

Equation (40) is explained as follows: The first term is simply the combined weight of paths of length 0. Of the remaining paths originating from A , we split into two cases depending on whether or not the path visits x . The second term combines the weights of paths that do not visit x . Furthermore, any path that does visit x can first spend $i \geq 0$ steps in B , and then $j \geq 0$ more steps in B after having visited x . Since the path has length at most m , we must have $i + j \leq m - 1$. This constitutes the third term in the summation. We may also write

$$LB_m(A + x) = 1 + |A| + \sum_{i=1}^m \omega(A \rightarrow B^i) + \sum_{i=1}^m \omega(x \rightarrow B^i).$$

We now define

$$\begin{aligned}
\Delta_{A,x} &:= LB_m(A+x) - LB_m(A) \\
&= 1 + \sum_{i=1}^m \omega(x \rightarrow B^i) - \sum_{\substack{i,j \geq 0 \\ i+j \leq m-1}} \omega(A \rightarrow B^i \rightarrow x \rightarrow B^j) \\
&= \left\{ 1 - \sum_{i=0}^{m-1} \omega(A \rightarrow B^i \rightarrow x) \right\} + \left\{ \sum_{i=1}^m \omega(x \rightarrow B^i) - \sum_{\substack{i \geq 0, j \geq 1 \\ i+j \leq m-1}} \omega(A \rightarrow B^i \rightarrow x \rightarrow B^j) \right\} \\
&:= I + II.
\end{aligned} \tag{41}$$

Note that $\sum_{i=0}^{m-1} \omega(A \rightarrow B^i \rightarrow x) \leq \omega(A \Rightarrow x)$. The latter is the probability of infecting x starting from the infection set A and is therefore bounded by 1, implying that $I \geq 0$. Furthermore,

$$\begin{aligned}
\sum_{\substack{i \geq 0, j \geq 1 \\ i+j \leq m-1}} \omega(A \rightarrow B^i \rightarrow x \rightarrow B^j) &= \sum_{j=1}^{m-1} \left(\sum_{i=0}^{m-1-j} \omega(A \rightarrow B^i \rightarrow x \rightarrow B^j) \right) \\
&\stackrel{(a)}{\leq} \sum_{j=1}^{m-1} \left(\sum_{i=0}^{m-1-j} \omega(A \rightarrow B^i \rightarrow x) \omega(x \rightarrow B^j) \right) \\
&= \sum_{j=1}^{m-1} \omega(x \rightarrow B^j) \left(\sum_{i=0}^{m-1-j} \omega(A \rightarrow B^i \rightarrow x) \right) \\
&\leq \sum_{j=1}^{m-1} \omega(x \rightarrow B^j) \omega(A \Rightarrow x) \\
&\leq \sum_{j=1}^{m-1} \omega(x \rightarrow B^j),
\end{aligned}$$

implying that $II \geq 0$. Note that inequality (a) holds because paths are not allowed to have redundant vertices. We will use a similar idea repeatedly in the proof of Lemma 3 below. Hence, we conclude that $\Delta_{A,x} \geq 0$, as wanted. \square

Next, we establish submodularity:

Lemma 3. *The function LB_m is submodular; i.e., for any disjoint sets A and B , and for any $x \notin A \cup B$, the following inequality holds:*

$$LB_m(A+x) - LB_m(A) \geq LB_m(A+B+x) - LB_m(A+B).$$

Proof. Using the notation in the proof of Lemma 2, we wish to show that $\Delta_{A+B,x} \leq \Delta_{A,x}$. Let $C := V \setminus (A+B+x)$. From the expressions in equation (41), we have

$$\Delta_{A+B,x} = 1 + \sum_{i=1}^m \omega(x \rightarrow C^i) - \sum_{\substack{i,j \geq 0 \\ i+j \leq m-1}} \omega((A+B) \rightarrow C^i \rightarrow x \rightarrow C^j).$$

Similarly, we may write

$$\Delta_{A,x} = 1 + \sum_{i=1}^m \omega(x \rightarrow (B+C)^i) - \sum_{\substack{i,j \geq 0 \\ i+j \leq m-1}} \omega(A \rightarrow (B+C)^i \rightarrow x \rightarrow (B+C)^j).$$

Letting

$$\mathcal{X} := \left(\sum_{i=1}^m \omega(x \rightarrow (B+C)^i) - \sum_{i=1}^m \omega(x \rightarrow C^i) \right) + \left(\sum_{\substack{i,j \geq 0 \\ i+j \leq m-1}} \omega(B \rightarrow C^i \rightarrow x \rightarrow C^j) \right)$$

and

$$\mathcal{Y} := \sum_{\substack{i,j \geq 0 \\ i+j \leq m-1}} \omega(A \rightarrow (B+C)^i \rightarrow x \rightarrow (B+C)^j) - \sum_{\substack{i,j \geq 0 \\ i+j \leq m-1}} \omega(A \rightarrow C^i \rightarrow x \rightarrow C^j),$$

the desired inequality is equivalent to $\mathcal{Y} \leq \mathcal{X}$.

Note that the first term in \mathcal{X} is the total weight of paths of length at most m that start from x and traverse between vertices in $B+C$, but visit at least one vertex in B . We denote the set of such paths by X_1 , and denote the set of paths appearing in the second term by X_2 . Then $\mathcal{X} = \omega(X_1) + \omega(X_2)$. We also introduce the notation

$$X_1 = \bigcup_{i=1}^m (x \rightarrow (B+C)^i)_{B \geq 1}. \quad (42)$$

On the other hand, the expression for \mathcal{Y} sums the weights of paths in

$$A \rightarrow (B+C)^i \rightarrow x \rightarrow (B+C)^j$$

such that at least one vertex in B is visited before or after visiting x . Let Y_1 denote the set of paths that visit B at least once after visiting x , and let Y_2 denote the set of paths that visit B at least once before visiting x , but do not visit B after visiting x . Then $\mathcal{Y} = \omega(Y_1) + \omega(Y_2)$. Using similar notation to equation (42), we write

$$\begin{aligned} Y_1 &= \bigcup_{\substack{i,j \geq 0 \\ i+j \leq m-1}} (A \rightarrow (B+C)^i \rightarrow x \rightarrow (B+C)^j)_{B \geq 1}, \quad \text{and} \\ Y_2 &= \bigcup_{\substack{i,j \geq 0 \\ i+j \leq m-1}} (A \rightarrow (B+C)^i)_{B \geq 1} \rightarrow x \rightarrow C^j. \end{aligned}$$

We then have the following sequence of inequalities:

$$\begin{aligned}
\omega(Y_1) &= \sum_{\substack{i+j \geq 0 \\ i+j \leq m-1}} \omega \left(A \rightarrow (B+C)^i \rightarrow x \rightarrow (B+C)^j_{B \geq 1} \right) \\
&= \sum_{j=1}^{m-1} \left(\sum_{i=0}^{m-1-j} \omega \left(A \rightarrow (B+C)^i \rightarrow x \rightarrow (B+C)^j_{B \geq 1} \right) \right) \\
&\leq \sum_{j=1}^{m-1} \left(\sum_{i=0}^{m-1-j} \omega \left(A \rightarrow (B+C)^i \rightarrow x \right) \omega \left(x \rightarrow (B+C)^j_{B \geq 1} \right) \right) \\
&= \sum_{j=1}^{m-1} \omega \left(x \rightarrow (B+C)^j_{B \geq 1} \right) \left(\sum_{i=0}^{m-1-j} \omega \left(A \rightarrow (B+C)^i \rightarrow x \right) \right) \\
&\leq \sum_{j=1}^{m-1} \omega \left(x \rightarrow (B+C)^j_{B \geq 1} \right) \omega(A \Rightarrow x) \\
&\leq \sum_{j=1}^{m-1} \omega \left(x \rightarrow (B+C)^j_{B \geq 1} \right) \\
&\leq \omega(X_1).
\end{aligned}$$

To analyze $\omega(Y_2)$, we partition the set of paths in Y_2 based on the last vertex visited in $B =$

$\{b_1, b_2, \dots, b_{|B|}\}$ before visiting x . Then

$$\begin{aligned}
\omega(Y_2) &= \sum_{\substack{i+j \geq 0 \\ i+j \leq m-1}} \omega(A \rightarrow (B+C)_{B \geq 1}^i \rightarrow x \rightarrow C^j) \\
&= \sum_{\substack{k_1, k_2, j \geq 0, \\ k_1+k_2+j \leq m-2}} \omega(A \rightarrow (B+C)^{k_1} \rightarrow B \rightarrow C^{k_2} \rightarrow x \rightarrow C^j) \\
&= \sum_{\substack{k_1, k_2, j \geq 0, \\ k_1+k_2+j \leq m-2 \\ 1 \leq \ell \leq |B|}} \omega(A \rightarrow (B+C)^{k_1} \rightarrow b_\ell \rightarrow C^{k_2} \rightarrow x \rightarrow C^j) \\
&\leq \sum_{\substack{k_1, k_2, j \geq 0, \\ k_1+k_2+j \leq m-2 \\ 1 \leq \ell \leq |B|}} \omega(A \rightarrow (B+C)^{k_1} \rightarrow b_\ell) \omega(b_\ell \rightarrow C^{k_2} \rightarrow x \rightarrow C^j) \\
&= \sum_{\substack{k_2, j \geq 0 \\ k_2+j \leq m-2 \\ 1 \leq \ell \leq |B|}} \omega(b_\ell \rightarrow C^{k_2} \rightarrow x \rightarrow C^j) \left(\sum_{k_1=0}^{m-2-j-k_2} \omega(A \rightarrow (B+C)^{k_1} \rightarrow b_\ell) \right) \\
&\leq \sum_{\substack{k_2, j \geq 0 \\ k_2+j \leq m-2 \\ 1 \leq \ell \leq |B|}} \omega(b_\ell \rightarrow C^{k_2} \rightarrow x \rightarrow C^j) \omega(A \Rightarrow b_\ell) \\
&\leq \sum_{\substack{k_2, j \geq 0 \\ k_2+j \leq m-2 \\ 1 \leq \ell \leq |B|}} \omega(b_\ell \rightarrow C^{k_2} \rightarrow x \rightarrow C^j) \\
&= \sum_{\substack{k_2, j \geq 0 \\ k_2+j \leq m-2}} \omega(B \rightarrow C^{k_2} \rightarrow x \rightarrow C^j) \\
&\leq \omega(X_2).
\end{aligned}$$

Thus, $\mathcal{Y} \leq \mathcal{X}$, completing the proof. \square

A.2.2 Proof of Theorem 8

The idea of the proof is relatively straightforward. For a given uninfected node i , adding vertices to the initial set of infected nodes may only increase the probability of the most probable infection path. This is more plainly evident in the proof of monotonicity, so we begin by establishing monotonicity. Let $S \subseteq T \subseteq V$. We need to show that

$$LB_{\text{trig}}(S) \leq LB_{\text{trig}}(T).$$

We start by writing

$$LB_{\text{trig}}(S) = \sum_{i \in V} \sup_{p \in P_{S,i}} \mathbb{E}[Y_p] = \sum_{i \in S} \sup_{p \in P_{S,i}} \mathbb{E}[Y_p] + \sum_{i \in \bar{S}} \sup_{p \in P_{S,i}} \mathbb{E}[Y_p].$$

Note that all the terms of the first sum are equal to 1. For any $i \in \bar{S}$, we claim that

$$\sup_{p \in P_{S,i}} \mathbb{E}[Y_p] \leq \sup_{p \in P_{T,i}} \mathbb{E}[Y_p]. \quad (43)$$

Indeed, suppose p is a path in $P_{S,i}$. If p does not pass through any vertex in $T \setminus S$, then p is also a path in $P_{T,i}$. If p does pass through a vertex in $T \setminus S$, then there is a path p' in $P_{T,i}$ that is contained within p . As a result, we see that

$$\mathbb{E}[Y_p] = \mathbb{P}\{Y_p = 1\} \leq \mathbb{P}\{Y_{p'} = 1\} = \mathbb{E}[Y_{p'}].$$

This proves equation (43), implying that

$$LB_{\text{trig}}(S) \leq \sum_{i \in S} \sup_{p \in P_{T,i}} \mathbb{E}[Y_p] + \sum_{i \in \bar{S}} \sup_{p \in P_{T,i}} \mathbb{E}[Y_p] = LB_{\text{trig}}(T),$$

which proves monotonicity.

We now consider submodularity. Let $S \subseteq T \subseteq V$, and let $v \in V \setminus T$. We need to show that

$$LB_{\text{trig}}(T \cup \{v\}) - LB_{\text{trig}}(T) \leq LB_{\text{trig}}(S \cup \{v\}) - LB_{\text{trig}}(S). \quad (44)$$

The main intuition for this proof remains the same, but with an added subtlety. Let $i \in V \setminus T$. By adding v to the set of initially infected vertices S , the most probable path from the initially infected set to i may only change to become more probable. However, if we let T be the set of initially infected vertices and add v to T , the most probable path from $T \cup \{v\}$ to i may start in $T \setminus S$, meaning that adding vertices to the initially infected set has diminishing returns. More precisely, we write

$$\begin{aligned} LB_{\text{trig}}(T \cup \{v\}) - LB_{\text{trig}}(T) &= \sum_{i \in V} \left(\sup_{p \in P_{T \cup \{v\},i}} \mathbb{E}[Y_p] - \sup_{p \in P_{T,i}} \mathbb{E}[Y_p] \right) \\ &= \sum_{i \in V_{T,v}} \left(\sup_{p \in P_{T \cup \{v\},i}} \mathbb{E}[Y_p] - \sup_{p \in P_{T,i}} \mathbb{E}[Y_p] \right), \end{aligned}$$

where $V_{T,v}$ is the set of vertices for which any path p in $P_{T \cup \{v\},i}$ that maximizes $\mathbb{E}[Y_p]$ starts at vertex v . Similarly, we have

$$LB_{\text{trig}}(S \cup \{v\}) - LB_{\text{trig}}(S) = \sum_{i \in V_{S,v}} \left(\sup_{p \in P_{S \cup \{v\},i}} \mathbb{E}[Y_p] - \sup_{p \in P_{S,i}} \mathbb{E}[Y_p] \right).$$

At this point, we make three claims, which together establish equation (44):

(i) For any i in $V_{T,v}$,

$$\sup_{p \in P_{T \cup \{v\},i}} \mathbb{E}[Y_p] \leq \sup_{p \in P_{S \cup \{v\},i}} \mathbb{E}[Y_p]. \quad (45)$$

(ii) We have the containment $V_{T,v} \subseteq V_{S,v}$.

(iii) We have the inequality

$$\sup_{p \in P_{S,i}} \mathbb{E}[Y_p] \leq \sup_{p \in P_{T,i}} \mathbb{E}[Y_p].$$

To establish claim (i), let p in $P_{C \cup \{v\},i}$ be a maximizer of $\mathbb{E}[Y_p]$. Then p starts at v and does not include an element of T . As a result, the path p is also in $P_{S \cup \{v\},i}$, establishing the first claim. Claim (ii) is immediate by noting that any vertex i for which any path p in $P_{T \cup \{v\},i}$ maximizing $\mathbb{E}[Y_p]$

starts at v is also a vertex in $P_{S \cup \{v\}, i}$. Note that claim (iii) is simply inequality (43). As a result of these three claims, we have

$$\begin{aligned} LB_{\text{trig}}(T \cup \{v\}) - LB_{\text{trig}}(T) &= \sum_{i \in V_{T,v}} \left(\sup_{p \in P_{T \cup \{v\}, i}} \mathbb{E}[Y_p] - \sup_{p \in P_{T,i}} \mathbb{E}[Y_p] \right) \\ &\leq \sum_{i \in V_{S,v}} \left(\sup_{p \in P_{S \cup \{v\}, i}} \mathbb{E}[Y_p] - \sup_{p \in P_{S,i}} \mathbb{E}[Y_p] \right) \\ &= LB_{\text{trig}}(S \cup \{v\}) - LB_{\text{trig}}(S), \end{aligned}$$

since each summand in the right hand side of the first and second lines is nonnegative. This completes the proof of submodularity.

B Additional Online Upper Bound Proofs

In this Appendix, we provide proofs for the pseudo-regret of player strategies based on the Exp3 and OSMD algorithms. We begin by proving Theorem 9, concerning the Exp3 algorithm, in Appendix B.1, and then prove Theorem 10, concerning the OSMD algorithm, in Appendix B.2.

B.1 Proof of Theorem 9

It suffices to verify the conditions of Proposition 3. We break down the analysis into the following series of lemmas:

Lemma 4. *We have the equalities*

$$\mathbb{E}_{S_t \sim p_t} \widehat{\ell}_{i,t}^{\text{sym}} = \ell_{i,t}, \quad \text{and} \quad \mathbb{E}_{I \sim p_t} \widehat{\ell}_{I,t}^{\text{sym}} = \ell_{S_t,t}.$$

Proof. The first claim follows from an easy calculation:

$$\begin{aligned} \mathbb{E}_{S_t \sim p_t} \widehat{\ell}_{i,t}^{\text{sym}} &= \frac{1}{n} \sum_{j \neq i} \frac{\ell_{i,j}^t}{p_{i,t} + p_{j,t}} \mathbb{E}_{S_t \sim p_t} [Z_{ij}] \\ &= \frac{1}{n} \sum_{j \neq i} \frac{\ell_{i,j}^t}{p_{i,t} + p_{j,t}} (p_{i,t} + p_{j,t}) \\ &= \frac{1}{n} \sum_{j \neq i} \ell_{i,j}^t \\ &= \ell_{i,t}. \end{aligned}$$

For the second claim, we write

$$\begin{aligned} \mathbb{E}_{I \sim p_t} \widehat{\ell}_{I,t}^{\text{sym}} &= \frac{1}{n} \sum_{i=1}^n p_{i,t} \sum_{j \neq i} \frac{\ell_{i,j}^t}{p_{i,t} + p_{j,t}} Z_{ij} \\ &= \frac{1}{2n} \sum_{i=1}^n \sum_{j \neq i} p_{i,t} \frac{\ell_{i,j}^t}{p_{i,t} + p_{j,t}} Z_{ij} + \frac{1}{2n} \sum_{i=1}^n \sum_{j \neq i} p_{j,t} \frac{\ell_{i,j}^t}{p_{i,t} + p_{j,t}} Z_{ij}, \end{aligned}$$

using symmetry in i and j to obtain the second equality. Hence,

$$\mathbb{E}_{I \sim p_t} \widehat{\ell}_{I,t}^{\text{sym}} = \frac{1}{2n} \sum_{i=1}^n \sum_{j \neq i} \ell_{i,j}^t Z_{ij} = \ell_{\mathcal{S}_t, t}.$$

Note that the factor of $\frac{1}{2}$ vanishes in the second inequality because we double-count each pair of vertices in the sum. \square

Lemma 5. *We have the inequality*

$$\mathbb{E}_{\mathcal{S} \sim p_t} \mathbb{E}_{I \sim p_t} (\widehat{\ell}_{i,t}^{\text{sym}})^2 \leq \frac{n+1}{2}.$$

Proof. Let $Z_i := \mathbf{1}_{\mathcal{S}_t = \{i\}}$. We have

$$\begin{aligned} \mathbb{E}_{I \sim p_t} (\widehat{\ell}_{i,t}^{\text{sym}})^2 &= \sum_{i=1}^n p_{i,t} \left(\frac{1}{n} \sum_{j \neq i} \ell_{i,j}^t \frac{Z_{ij}}{p_{i,t} + p_{j,t}} \right)^2 \\ &= \sum_{i=1}^n \frac{p_{i,t}}{n^2} \sum_{j \neq i} \sum_{k \neq i} \ell_{i,j}^t \ell_{i,k}^t \frac{Z_{ij} Z_{ik}}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{k,t})} \\ &= \sum_{i=1}^n \frac{p_{i,t}}{n^2} \sum_{j \neq i} \sum_{k \neq i} \ell_{i,j}^t \ell_{i,k}^t \frac{Z_{ij} Z_{ik}}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{k,t})} Z_i \\ &\quad + \sum_{i=1}^n \frac{p_{i,t}}{n^2} \sum_{j \neq i} \ell_{i,j}^t \ell_{i,j}^t \frac{Z_{ij} Z_{ij}}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{j,t})} Z_j \\ &:= E_1 + E_2. \end{aligned}$$

Note that the last equality follows from the observation that for any summand to be nonzero, both Z_{ij} and Z_{ik} must be nonzero. As a result, we only sum the terms for which $\mathcal{S}_t = \{i\}$ or $\mathcal{S}_t = \{j\} = \{k\}$.

Now, we calculate the expectation of E_1 . We may write

$$\begin{aligned} E_1 &= \frac{p_{\mathcal{S}_t, t}}{n^2} \sum_{j \neq \mathcal{S}_t} \sum_{k \neq \mathcal{S}_t} \frac{\ell_{\mathcal{S}_t, j}^t \ell_{\mathcal{S}_t, k}^t}{(p_{\mathcal{S}_t, t} + p_{j,t})(p_{\mathcal{S}_t, t} + p_{k,t})} \\ &= \left(\frac{1}{n} \sum_{j \neq \mathcal{S}_t} \frac{p_{\mathcal{S}_t, t}}{p_{\mathcal{S}_t, t} + p_{j,t}} \ell_{\mathcal{S}_t, j}^t \right) \left(\frac{1}{n} \sum_{k \neq \mathcal{S}_t} \frac{\ell_{\mathcal{S}_t, k}^t}{p_{\mathcal{S}_t, t} + p_{k,t}} \right) \\ &\leq \left(\frac{1}{n} \sum_{k \neq \mathcal{S}_t} \frac{\ell_{\mathcal{S}_t, k}^t}{p_{\mathcal{S}_t, t} + p_{k,t}} \right), \end{aligned}$$

where the first equality holds because only one of the terms in the sum over i is nonzero. Finally, note that

$$\mathbb{E}_{\mathcal{S}_t \sim p_t} \left[\frac{1}{n} \sum_{k \neq \mathcal{S}_t} \frac{\ell_{\mathcal{S}_t, k}^t}{p_{\mathcal{S}_t, t} + p_{k,t}} \right] = \frac{1}{n} \sum_{i=1}^n \sum_{k \neq i} \frac{p_{i,t}}{p_{i,t} + p_{k,t}} \ell_{i,j}^t \leq \frac{1}{n} \sum_{i=1}^n \sum_{k \neq i} \frac{p_{i,t}}{p_{i,t} + p_{k,t}}.$$

Furthermore, we have the useful equation

$$\sum_{i=1}^n \sum_{k=1}^n \frac{a_i}{a_i + a_k} = \frac{n^2}{2}, \quad (46)$$

for any nonnegative sequence $\{a_i\}_{i=1}^n$. This may be seen via the following algebraic manipulations:

$$\begin{aligned} \sum_{i=1}^n \sum_{k=1}^n \frac{a_i}{a_i + a_k} &= \sum_{i=1}^n \frac{a_i}{a_i + a_i} + \sum_{k \neq i} \frac{a_i}{a_i + a_k} \\ &= \frac{n}{2} + \frac{1}{2} \sum_{k \neq i} \frac{a_i}{a_i + a_k} + \frac{1}{2} \sum_{k \neq i} \frac{a_k}{a_i + a_k} \\ &= \frac{n}{2} + \frac{1}{2} \sum_{k \neq i} \frac{a_i + a_k}{a_i + a_k} \\ &= \frac{n}{2} + \frac{n(n-1)}{2} \\ &= \frac{n^2}{2}. \end{aligned}$$

At this point, we calculate the expectation of E_2 . We first obtain

$$E_2 = \sum_{i=1}^n \frac{p_{i,t}}{n^2} \sum_{j \neq i} \frac{\ell_{i,j}^t}{(p_{i,t} + p_{j,t})^2} Z_j \leq \sum_{i=1}^n \frac{p_{i,t}}{n^2} \sum_{j \neq i} \frac{Z_j}{(p_{i,t} + p_{j,t})^2}.$$

Taking $\mathbb{E}_{\mathcal{S}_t \sim p_t}$, we obtain

$$\mathbb{E}_{\mathcal{S} \sim p_t} E_2 \leq \frac{1}{n^2} \sum_{i \neq j} \frac{p_{i,j} p_{j,t}}{(p_{i,t} + p_{j,t})^2} \stackrel{(a)}{\leq} \frac{1}{n^2} \frac{n^2}{2} = \frac{1}{2},$$

where in (a), we have used the fact that $\frac{ab}{(a+b)^2} \leq \frac{1}{2}$ for all $a, b \in \mathbb{R}$.

Altogether, we have

$$\mathbb{E}_{\mathcal{S} \sim p_t} \mathbb{E}_{I \sim p_t} (\widehat{\ell}_{i,t}^{\text{sym}})^2 \leq \mathbb{E}_{\mathcal{S}_t \sim p_t} \left[\frac{1}{n} \sum_{k \neq \mathcal{S}_t} \frac{\ell_{\mathcal{S}_t,k}^t}{p_{\mathcal{S}_t,t} + p_{k,t}} \right] + \frac{1}{2} \leq \frac{n}{2} + \frac{1}{2} = \frac{n+1}{2},$$

completing the proof. \square

This also completes the proof of Theorem 9.

B.2 OSMD Proofs

The goal of this Appendix is to prove Theorem 10. We begin with some preliminaries.

B.2.1 Preliminaries

We first describe the function F_ψ . Recall that a continuous function $F : \overline{\mathcal{D}} \rightarrow \mathbb{R}$ is a Legendre function if F is strictly convex, F has continuous first partial derivatives on \mathcal{D} , and

$$\lim_{x \rightarrow \overline{\mathcal{D}} \setminus \mathcal{D}} \|\nabla F(x)\| = \infty.$$

The analysis in this paper concerns a very specific type of Legendre function associated to a 0-potential, as described in the following definition:

Definition 3. A function $\psi : (-\infty, a) \rightarrow \mathbb{R}_+$ is called a 0-potential if it is convex, continuously differentiable, and satisfies the following conditions:

$$\begin{aligned} \lim_{x \rightarrow -\infty} \psi(x) &= 0, & \lim_{x \rightarrow a} \psi(x) &= \infty, \\ \psi' &> 0, & \int_0^1 |\psi^{-1}(s)| ds &\leq \infty. \end{aligned}$$

We additionally define the associated function F_ψ on $(0, \infty)^n$ by

$$F_\psi(x) = \sum_{i=1}^n \int_0^{x_i} \psi^{-1}(s) ds.$$

In particular, we will consider the 0-potential $\psi(x) = (-x)^{-q}$. Then $\psi^{-1}(x) = -x^{-\frac{1}{q}}$, so

$$F_\psi(x) = -\frac{q}{q-1} \sum_{i=1}^n x_i^{\frac{q-1}{q}}.$$

Specifically, we will consider the case $q = 2$ (the same analysis could be performed with respect to $q > 1$, and then the final bound could be optimized over q).

To employ Proposition 4, we need to bound two summands. The following simple lemma bounds the first term:

Lemma 6. When $\psi(x) = \frac{1}{x^2}$, we have the bound

$$F_\psi(p) - F_\psi(p_1) \leq 2\sqrt{n}, \quad \forall p \in \Delta^n.$$

Proof. Since $F_\psi(p) \leq 0$ and $\|p_1\|_1 = 1$, Hölder's inequality implies that

$$F_\psi(p) - F_\psi(p_1) \leq 2 \sum_{i=1}^n p_{1,i}^{1/2} \leq 2n^{\frac{1}{2}}.$$

This completes the proof of the lemma. □

All that remains is to analyze the loss-specific term appearing in Proposition 4 and choose η appropriately.

B.2.2 Proof of Theorem 10

We first prove the following lemma:

Lemma 7. We have the inequality

$$\sum_{i=1}^n \mathbb{E} \left[\frac{(\widehat{\ell}_{i,t}^{sym})^2}{(\psi^{-1})'(p_{i,t})} \right] \leq \sqrt{2n}, \quad \forall 1 \leq t \leq T. \quad (47)$$

Proof. Let \mathcal{F}_t denote the sigma-field of all actions up to time t . We have

$$\begin{aligned}
\sum_{i=1}^n \mathbb{E} \left[\frac{(\widehat{\ell}_{i,t}^{\text{sym}})^2}{(\psi^{-1})'(p_{i,t})} \middle| \mathcal{F}_{t-1} \right] &\stackrel{(a)}{=} 2 \sum_{i=1}^n p_{i,t}^{3/2} \mathbb{E} \left[(\widehat{\ell}_{i,t}^{\text{sym}})^2 \middle| \mathcal{F}_{t-1} \right] \\
&\stackrel{(b)}{\leq} 2 \left(\sum_{i=1}^n p_{i,t} \right)^{1/2} \left(\sum_{i=1}^n \left(p_{i,t} \mathbb{E} \left[(\widehat{\ell}_{i,t}^{\text{sym}})^2 \middle| \mathcal{F}_{t-1} \right] \right)^2 \right)^{1/2} \\
&= 2 \left(\sum_{i=1}^n \left(p_{i,t} \mathbb{E} \left[(\widehat{\ell}_{i,t}^{\text{sym}})^2 \middle| \mathcal{F}_{t-1} \right] \right)^2 \right)^{1/2}, \tag{48}
\end{aligned}$$

where we have used the facts that $(\psi^{-1})'(x) = \frac{1}{2}x^{-3/2}$ and p_t is measurable with respect to \mathcal{F}_{t-1} to establish (a), and applied Hölder's inequality to obtain (b).

We now inspect the conditional expectation more closely. We have

$$\begin{aligned}
\mathbb{E} \left[(\widehat{\ell}_{i,t}^{\text{sym}})^2 \middle| \mathcal{F}_{t-1} \right] &= \mathbb{E} \left[\left(\frac{1}{n} \sum_{j \neq i} \frac{1}{p_{i,t} + p_{j,t}} \ell_{i,j}^t Z_{ij} \right)^2 \middle| \mathcal{F}_{t-1} \right], \\
&= \frac{1}{n^2} \mathbb{E} \left[\sum_{j \neq i} \sum_{k \neq i} \frac{1}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{k,t})} \ell_{i,j}^t \ell_{i,k}^t Z_{ij} Z_{ik} \middle| \mathcal{F}_{t-1} \right] \\
&= \frac{1}{n^2} \mathbb{E} \left[\sum_{j \neq i} \sum_{k \neq i} \frac{1}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{k,t})} \ell_{i,j}^t \ell_{i,k}^t Z_i \middle| \mathcal{F}_{t-1} \right] \\
&\quad + \frac{1}{n^2} \mathbb{E} \left[\sum_{j \neq i} \frac{1}{(p_{i,t} + p_{j,t})^2} (\ell_{i,j}^t)^2 Z_j \middle| \mathcal{F}_{t-1} \right],
\end{aligned}$$

where the third equality is due to the fact that $Z_{ij}Z_{ik}$ is 1 only when i is the source vertex or $j = k$ is the source vertex. Using the fact that $\ell_{i,j}^t$ is bounded by 1, we then obtain

$$\begin{aligned}
\mathbb{E} \left[(\widehat{\ell}_{i,t}^{\text{sym}})^2 \middle| \mathcal{F}_{t-1} \right] &\leq \frac{1}{n^2} \mathbb{E} \left[\sum_{j \neq i} \sum_{k \neq i} \frac{Z_i}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{k,t})} \middle| \mathcal{F}_{t-1} \right] \\
&\quad + \frac{1}{n^2} \mathbb{E} \left[\sum_{j \neq i} \frac{Z_j}{(p_{i,t} + p_{j,t})^2} \middle| \mathcal{F}_{t-1} \right] \\
&\leq \frac{1}{n^2} \sum_{j \neq i} \sum_{k \neq i} \frac{p_{i,t}}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{k,t})} + \frac{1}{n^2} \sum_{j \neq i} \frac{p_{j,t}}{(p_{i,t} + p_{j,t})^2} \\
&\leq \frac{1}{n^2} \sum_{j \neq i} \sum_{k \neq i} \frac{p_{i,t} + p_{j,t}}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{k,t})} \\
&= \frac{1}{n^2} \sum_{j \neq i} \sum_{k \neq i} \frac{1}{p_{i,t} + p_{k,t}} \\
&\leq \frac{1}{n} \sum_{k=1}^n \frac{1}{p_{i,t} + p_{k,t}}.
\end{aligned}$$

Combining this result with the bound (48), we have

$$\begin{aligned}
\sum_{i=1}^n \mathbb{E} \left[\frac{(\widehat{\ell}_{i,t}^{\text{sym}})^2}{(\psi^{-1})'(p_{i,t})} \middle| \mathcal{F}_{t-1} \right] &\leq 2 \left(\sum_{i=1}^n \left(\frac{p_{i,t}}{n} \sum_{k=1}^n \frac{1}{p_{i,t} + p_{k,t}} \right)^2 \right)^{1/2} \\
&= \frac{2}{n} \left(\sum_{i=1}^n \sum_{j=1}^n \sum_{k=1}^n \frac{p_{i,t}^2}{(p_{i,t} + p_{j,t})(p_{i,t} + p_{k,t})} \right)^{\frac{1}{2}} \\
&\leq \frac{2}{n} \left(n \sum_{i=1}^n \sum_{j=1}^n \frac{p_{i,t}}{p_{i,t} + p_{j,t}} \right)^{\frac{1}{2}}.
\end{aligned}$$

Appealing to equation (46), we may replace the double sum by $\frac{n^2}{2}$ and simplify the bound:

$$\sum_{i=1}^n \mathbb{E} \left[\frac{(\widehat{\ell}_{i,t}^{\text{sym}})^2}{(\psi^{-1})'(p_{i,t})} \middle| \mathcal{F}_{t-1} \right] \leq \frac{2}{n} \left(\frac{n^3}{2} \right)^{\frac{1}{2}} = \sqrt{2n}.$$

Taking an additional expectation and using the tower property, we arrive at the desired inequality. \square

Combining Lemmas 6 and 7 with Proposition 4, we then have

$$\sup_{\mathcal{A} \in \mathcal{A}} \overline{R}_T(\mathcal{A}, \mathcal{S}_{\text{OSMD}}^{\text{sym}}) \leq \frac{2\sqrt{n}}{\eta} + \eta T \sqrt{\frac{n}{2}}.$$

Optimizing over η , we take $\eta = 2^{\frac{3}{4}} T^{-\frac{1}{2}}$, which establishes the desired bound.

C Adversarial Influence Maximization with Multiple Sources

In this Appendix, we prove results concerning multiple infection sources. Fix an adversarial strategy \mathcal{A} , and define the functions $f_t(\mathcal{S}_t) = f(\mathcal{A}_t, \mathcal{S}_t)$ and $F(\mathcal{S}) = \sum_{t=1}^T f_t(\mathcal{S}_t)$. Thus, $F(\mathcal{S})$ is the total reward for strategy $\mathcal{S} = \{\mathcal{S}_t\}$. In the stochastic setting, when $T = 1$, many influence-maximization analyses exploit the submodularity of f_t under certain stochastic assumptions on \mathcal{A}_t . In the bandit setting, we wish to establish an analogous result for F , in order to establish regret bounds when the player chooses source vertices according to a greedy algorithm.

Since $\mathcal{S} \in (2^V)^T$, the function F is not technically a set function. However, we may identify each player strategy \mathcal{S} with an element of $\mathcal{S}^* \in 2^{(V^T)}$, and define $F^*(\mathcal{S}^*) = F(\mathcal{S})$. Here,

$$V^T := \{v^T = (v(1), v(2), \dots, v(T)) \mid v(i) \in V, \text{ for } 1 \leq i \leq T\},$$

and $\mathcal{S}^* = \{u_1^T, \dots, u_r^T\} \in 2^{(V^T)}$ corresponds to the strategy that selects the source nodes $\{u_1(t), \dots, u_r(t)\}$ at time t .

We first show that F^* is a monotone, submodular function:

Lemma 8. *The function $f_t(\mathcal{S}_t) = f(\mathcal{A}_t, \mathcal{S}_t)$ is monotone and submodular, for every fixed \mathcal{A}_t .*

Proof. It is trivial to see that f_t is monotone, so we focus on proving submodularity. Our goal is to show that for a fixed \mathcal{A}_t , and for any $\mathcal{S}_t \subseteq \mathcal{S}'_t$ and $u \in V \setminus \mathcal{S}'_t$, we have

$$f_t(\mathcal{S}'_t \cup \{u\}) - f_t(\mathcal{S}'_t) \leq f_t(\mathcal{S}_t \cup \{u\}) - f_t(\mathcal{S}_t). \tag{49}$$

Let $Z_{\mathcal{S}_t, v}$ denote the indicator of an open path between a source node $s \in \mathcal{S}_t$ and $v \in V$, where by convention, $Z_{\mathcal{S}_t, v} = 1$ if $v \in \mathcal{S}_t$. Note that $f_t(\mathcal{S}_t) = \frac{1}{n} \sum_{v \in V} Z_{\mathcal{S}_t, v}$. We will show that for $v \notin \mathcal{S}'_t \cup \{u\}$, we have

$$Z_{\mathcal{S}'_t \cup \{u\}, v} - Z_{\mathcal{S}'_t, v} \leq Z_{\mathcal{S}_t \cup \{u\}, v} - Z_{\mathcal{S}_t, v}. \quad (50)$$

Summing over $v \in (\mathcal{S}'_t \cup \{u\})^c$, using $Z_{\mathcal{S}_t \cup \{u\}, v} - Z_{\mathcal{S}_t, v} \geq 0$ for $v \in \mathcal{S}'_t \setminus \mathcal{S}_t$, and dividing by n will yield the desired inequality (49).

We have three cases to consider: In the first case, an open path exists from some $s \in \mathcal{S}'_t$ to v . Then the left side of inequality (50) is equal to 0, while the right hand side is at least 0 by monotonicity. In the second case, an open path does not exist from any $s \in \mathcal{S}'_t$ to v , but an open path exists from u to v . Then both sides of inequality (50) are equal to 1. Finally, if no open path exists from $s \in \mathcal{S}'_t \cup \{u\}$ to v , then both sides of inequality (50) are equal to 0. This completes the proof. \square

Proposition 6. *The function F^* is monotone and submodular.*

Proof. The properties are essentially immediate from Lemma 8. Let \mathcal{P} and \mathcal{Q} be elements of $(2^V)^T$ such that $\mathcal{P}^* \subseteq \mathcal{Q}^*$. Then

$$F^*(\mathcal{P}^*) = \sum_{t=1}^T f_t(\mathcal{P}_t) \leq \sum_{t=1}^T f_t(\mathcal{Q}_t) = F^*(\mathcal{Q}^*),$$

proving monotonicity. Similarly, if $\mathcal{S} \in (2^V)^T$, we have

$$\begin{aligned} F^*(\mathcal{S}^* \cup \mathcal{Q}^*) - F^*(\mathcal{Q}^*) &= \sum_{t=1}^T (f_t(\mathcal{S}_t \cup \mathcal{Q}_t) - f_t(\mathcal{Q}_t)) \\ &\leq \sum_{t=1}^T (f_t(\mathcal{S}_t \cup \mathcal{P}_t) - f_t(\mathcal{P}_t)) \\ &= F^*(\mathcal{S}^* \cup \mathcal{P}^*) - F^*(\mathcal{P}^*), \end{aligned}$$

proving submodularity. \square

By Proposition 2, we then have

$$\left(1 - \frac{1}{e}\right) \max_{|\mathcal{S}^*| \leq K} F^*(\mathcal{S}^*) \leq F^*(G^*),$$

where G^* is a set of cardinality $K \geq 1$ constructed via a sequential greedy algorithm. However, this result is not immediately applicable to the online bandit setting, since we do not have direct access to F^* . Thus, we can only hope to obtain an approximate greedy maximizer \tilde{G}^* , and we wish to derive theoretical guarantees for $F^*(\tilde{G}^*)$.

Our result relies on the following general proposition:

Proposition 7 (Theorem 6 from [Streeter and Golovin, 2007](#)). *Let $f : 2^{\mathcal{V}} \rightarrow \mathbb{R}$ be a monotone, submodular function such that $f(\emptyset) = 0$. Consider a set $\mathcal{D} \subseteq \mathcal{V}$ and a sequence of error tolerances $\{\epsilon_i\}$, and suppose $\{G_i^\epsilon\}$ is constructed in an approximate greedy manner, such that $G_0^\epsilon = \emptyset$ and $G_i^\epsilon = G_{i-1}^\epsilon \cup \{g_i\}$, where*

$$\max_{d \in \mathcal{D}} f(G_{i-1}^\epsilon \cup \{d\}) - f(G_{i-1}^\epsilon) \leq f(G_{i-1}^\epsilon \cup \{g_i\}) - f(G_{i-1}^\epsilon) + \epsilon_i.$$

Then for any $K \geq 1$, we have

$$\left(1 - \frac{1}{e}\right) \max_{\mathcal{S}^* \in \mathcal{D}_K} f(\mathcal{S}^*) - f(G_K^\epsilon) \leq \sum_{i=1}^K \epsilon_i,$$

where \mathcal{D}_K consists of subsets of \mathcal{D} containing at most K elements.

Proposition 7 ensures that for submodular functions, successive errors $\{\epsilon_i\}$ in a sequential greedy algorithm only accumulate additively. The proof is provided in Streeter and Golovin (2007), but we include a proof in Appendix C.2 for completeness.

C.1 Proof of Theorem 11

Suppose $\mathcal{A} \in \mathcal{A}$. We will apply Proposition 7 with $f = \mathbb{E}_{\mathcal{A}} F^*$, $\mathcal{V} = V^T$, and $K = k$. Note that $\mathbb{E}_{\mathcal{A}} F^*$ inherits monotonicity and submodularity from F^* . Also let

$$\mathcal{D} = \{(v, \dots, v) : v \in V\} \subseteq 2^{(V^T)}$$

denote the diagonal set of $2^{(V^T)}$. For a (non-random) k -source strategy \mathcal{S}^* with $|\mathcal{S}_t^*| = k$ for all t , we use the notation $\mathcal{S}^* = \{\mathcal{S}_1^*, \dots, \mathcal{S}_k^*\}$, where \mathcal{S}_i^* corresponds to the set of i^{th} vertices chosen during the T rounds. Proposition 7 immediately gives

$$\left(1 - \frac{1}{e}\right) \max_{\mathcal{S}^* \in \mathcal{D}} \mathbb{E}_{\mathcal{A}} F^*(\mathcal{S}^*) - \mathbb{E}_{\mathcal{A}} F^*(G_k^\epsilon) \leq \sum_{i=1}^k \max_{d_i \in \mathcal{D}} \mathbb{E}_{\mathcal{A}} [F^*(G_{i-1}^\epsilon \cup \{d_i\}) - F^*(G_{i-1}^\epsilon \cup \{g_i\})],$$

where the sets $\{G_i^\epsilon\}$ are chosen in an approximate greedy manner, and we have taken ϵ_i to be minimal. In particular, we consider $\{G_i^\epsilon\}$ to be the choice of i^{th} vertices \mathcal{S}_i^* corresponding to the player's choice under the strategy \mathcal{S}^1 .

We now take an expectation with respect to possible randomization in the player's strategy, to obtain

$$\begin{aligned} \overline{R}_T^{(1-1/e)}(\mathcal{A}, \mathcal{S}) &\leq \sum_{i=1}^k \mathbb{E}_{\mathcal{S}} \left[\max_{d_i \in \mathcal{D}} \mathbb{E}_{\mathcal{A}} [F^*(G_{i-1}^\epsilon \cup \{d_i\}) - F^*(G_{i-1}^\epsilon \cup \{g_i\})] \right] \\ &\stackrel{(a)}{=} \sum_{i=1}^k \mathbb{E}_{\mathcal{S}_{[1:i]}} \left[\max_{d_i \in \mathcal{D}} \mathbb{E}_{\mathcal{A}} [F^*(G_{i-1}^\epsilon \cup \{d_i\}) - F^*(G_{i-1}^\epsilon \cup \{g_i\})] \right] \\ &\stackrel{(b)}{=} \sum_{i=1}^k \mathbb{E}_{\mathcal{S}_{[1:i-1]}} \left[\mathbb{E}_{\mathcal{S}_i} \max_{d_i \in \mathcal{D}} \mathbb{E}_{\mathcal{A}} [F^*(G_{i-1}^\epsilon \cup \{d_i\}) - F^*(G_{i-1}^\epsilon \cup \{g_i\})] \right]. \end{aligned}$$

Here, $\mathbb{E}_{\mathcal{S}_{[1:i]}}$ denotes the expectation with respect to the first i vertices played, and the equality in (a) holds because the set of i^{th} vertices played depends only on the sets of the first i vertices played. The equality in (b) holds because the set G_{i-1}^ϵ , and hence the choice of d_i , does not depend on the selection of i^{th} vertices. Furthermore, the inner expression is simply the pseudo-regret of strategy \mathcal{S}^1 . By Theorem 10, this is bounded by $2^{\frac{1}{4}} \sqrt{nT}$. Summing up, we obtain the desired result.

C.2 Proof of Proposition 7

We begin with two supporting lemmas:

Lemma 9. For any $\mathcal{P} \subseteq \mathcal{V}$ and $\mathcal{Q} \subseteq \mathcal{D}$, we have

$$f(\mathcal{P} \cup \mathcal{Q}) \leq f(\mathcal{P}) + |\mathcal{Q}| \max_{v \in \mathcal{D}} [f(\mathcal{P} \cup \{v\}) - f(\mathcal{P})].$$

Proof. We proceed by induction on $|\mathcal{Q}|$. The case $|\mathcal{Q}| = 1$ is immediate. Now suppose the statement is true for all $|\mathcal{Q}| \leq k$, where $k \geq 1$. Let $c \in \mathcal{D}$, and suppose $\mathcal{Q} \subseteq \mathcal{D}$ has cardinality k . Then

$$\begin{aligned} f(\mathcal{P} \cup (\mathcal{Q} \cup \{c\})) &\stackrel{(a)}{\leq} f(\mathcal{P} \cup \{c\}) + |\mathcal{Q}| \max_{d \in \mathcal{D}} [f((\mathcal{P} \cup \{c\}) \cup \{d\}) - f(\mathcal{P} \cup \{c\})] \\ &\stackrel{(b)}{\leq} f(\mathcal{P}) + \max_{d \in \mathcal{D}} [f(\mathcal{P} \cup \{d\}) - f(\mathcal{P})] + |\mathcal{Q}| \max_{d \in \mathcal{D}} [f(\mathcal{P} \cup \{d\}) - f(\mathcal{P})] \\ &= f(\mathcal{P}) + |\mathcal{Q} \cup \{c\}| \max_{d \in \mathcal{D}} [f(\mathcal{P} \cup \{d\}) - f(\mathcal{P})], \end{aligned}$$

where (a) follows from the induction hypothesis and (b) follows from the induction hypothesis and submodularity. This completes the induction and proves the lemma. \square

Lemma 10. Let $\delta_i := f(G_i^\epsilon) - f(G_{i-1}^\epsilon)$. For any $\mathcal{Q} \subseteq \mathcal{D}$, we have

$$f(\mathcal{Q}) \leq f(G_{i-1}^\epsilon) + |\mathcal{Q}|(\delta_i + \epsilon_i).$$

Proof. Using Lemma 9 and monotonicity of f , we have

$$\begin{aligned} f(\mathcal{Q}) &\leq f(G_{i-1}^\epsilon \cup \mathcal{Q}) \\ &\leq f(G_{i-1}^\epsilon) + |\mathcal{Q}| \max_{d \in \mathcal{D}} [f(G_{i-1}^\epsilon \cup \{d\}) - f(G_{i-1}^\epsilon)] \\ &\leq f(G_{i-1}^\epsilon) + |\mathcal{Q}| (f(G_i^\epsilon) - f(G_{i-1}^\epsilon) + \epsilon_i) \\ &= f(G_{i-1}^\epsilon) + |\mathcal{Q}|(\delta_i + \epsilon_i), \end{aligned}$$

completing the proof. \square

We now define $\Delta_i := \max_{\mathcal{S}^* \in \mathcal{D}_K} f(\mathcal{S}^*) - f(G_{i-1}^\epsilon)$. By Lemma 10, we have

$$\max_{\mathcal{S}^* \in \mathcal{D}_K} f(\mathcal{S}^*) \leq f(G_{i-1}^\epsilon) + K(\delta_i + \epsilon_i).$$

Subtracting $f(G_{i-1}^\epsilon)$, we obtain

$$\Delta_i \leq K(\delta_i + \epsilon_i) = K(\Delta_i - \Delta_{i+1} + \epsilon_i),$$

so

$$\Delta_{i+1} \leq \Delta_i \left(1 - \frac{1}{K}\right) + \epsilon_i.$$

Applying this inequality recursively, we see that

$$\Delta_{K+1} \leq \Delta_1 \prod_{i=1}^K \left(1 - \frac{1}{K}\right) + \sum_{i=1}^K \epsilon_i = \Delta_1 \left(1 - \frac{1}{K}\right)^K + \sum_{i=1}^K \epsilon_i \leq \Delta_1 \left(\frac{1}{e}\right) + \sum_{i=1}^K \epsilon_i.$$

Rearranging and using the fact that $f(\emptyset) = 0$ completes the proof.

D Additional Online Lower Bound Proofs

The main goal of this Appendix is to prove Theorems 12 and 14. Some of the computations are rather lengthy and are therefore included in Appendix D.2.

D.1 Proofs of Theorems

We first present the main components of the proofs, followed by detailed calculations involving the Kullback-Leibler divergence.

D.1.1 Proof of Theorem 12

Let the adversarial strategies $\{\mathcal{A}^i\}$ be defined as follows: For each strategy, the adversary chooses a random subset of vertices, and opens all edges between vertices in the subset. For \mathcal{A}^i , with $1 \leq i \leq n$, the adversary includes vertex i with probability $\frac{c}{n}$, and includes all other vertices with probability $\frac{c}{n}(1 - \delta)$ each, where $\delta \in (0, 1/2)$ is a small constant. Finally, for \mathcal{A}^0 , the adversary includes all vertices independently with probability $\frac{c}{n}(1 - \delta)$. Successive actions of the adversary are i.i.d. across time steps.

We now derive the following lemmas, which will be used in Proposition 5:

Lemma 11. *For any $i \neq j$ and $1 \leq t \leq T$, we have*

$$\mathbb{E}_i[X_{i,t} - X_{j,t}] = \frac{(n-2)c^2}{n^3}(1-\delta)\delta.$$

Proof. Let \mathcal{C}_t be the clique chosen by the adversary at time t . Note that if $i, j \in \mathcal{C}_t$ or $i, j \notin \mathcal{C}_t$, the difference in rewards is 0. Thus, the only cases of interest in computing the expectation are when exactly one of i or j is in \mathcal{C}_t . Then

$$\begin{aligned} \mathbb{E}_i[X_{i,t} - X_{j,t}] &= \mathbb{E}_i[(|\mathcal{C}_t| - 1)\mathbf{1}_{i \in \mathcal{C}_t}\mathbf{1}_{j \notin \mathcal{C}_t} - (1 - |\mathcal{C}_t|)\mathbf{1}_{i \notin \mathcal{C}_t}\mathbf{1}_{j \in \mathcal{C}_t}] \\ &= \frac{1}{n} \left(\frac{c}{n}(n-2)(1-\delta) \right) \left(\frac{c}{n} \right) \left[1 - \frac{c}{n}(1-\delta) \right] \\ &\quad - \frac{1}{n} \left(\frac{c}{n}(n-2)(1-\delta) \right) \left[1 - \frac{c}{n} \right] \left[\frac{c}{n}(1-\delta) \right] \\ &= \frac{1}{n}(n-2) \left(\frac{c}{n} \right)^2 (1-\delta) \left(\left[1 - \frac{c}{n}(1-\delta) \right] - \left[1 - \frac{c}{n} \right] (1-\delta) \right) \\ &= \frac{(n-2)c^2}{n^3}(1-\delta)\delta, \end{aligned}$$

where the second equality uses the fact that $\frac{c}{n}(n-2)(1-\delta)$ other vertices are expected to be in \mathcal{C}_t . \square

Lemma 12. *Let $\mathcal{S} \in \mathcal{P}_d$ be a deterministic player strategy, and let $T_i = |\{t : \mathcal{S}_t = \{i\}\}|$. Then we have the upper bound*

$$\sum_{i=1}^n KL(\mathbb{P}_0, \mathbb{P}_i) \leq \frac{c(c+1)}{n-c} T \delta^2.$$

The proof of Lemma 12 is provided in Appendix D.2.1.

Thus, by Proposition 5, we have

$$\begin{aligned} \inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}) &\geq T \frac{(n-2)c^2}{n^3} (1-\delta) \delta \left(\frac{n-1}{n} - \delta \sqrt{\frac{T}{2n}} \sqrt{\frac{c}{n-c}} (c+1) \right) \\ &\geq \frac{T}{6} \left(\frac{c}{n} \right)^2 \left(\frac{n-1}{n} \delta - \delta^2 \sqrt{\frac{T}{2n}} \sqrt{\frac{c(c+1)}{n-c}} \right), \end{aligned}$$

where the second inequality uses the fact that $n \geq 3$ and $\delta < 1/2$. Finally, we optimize over δ and c . Since we have a quadratic equation in δ , we take

$$\delta = \frac{n-1}{2n} \sqrt{\frac{2n}{T}} \sqrt{\frac{n-c}{c(c+1)}},$$

yielding

$$\begin{aligned} \inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}) &\geq \frac{T}{6} \left(\frac{c}{n} \right)^2 \left(\frac{1}{4} \left(\frac{n-1}{n} \right)^2 \sqrt{\frac{2n}{T}} \sqrt{\frac{n-c}{c(c+1)}} \right) \\ &= \frac{1}{12\sqrt{2}} \sqrt{T} \left(\frac{c}{n} \right)^2 \left(\frac{n-1}{n} \right)^2 \sqrt{\frac{n(n-c)}{c(c+1)}} \\ &\geq \frac{1}{27\sqrt{3}} \sqrt{T} \left(\frac{c}{n^2} \right) \sqrt{n(n-c)}, \end{aligned}$$

where the second inequality uses the bounds $\frac{n-1}{n} \geq \frac{2}{3}$ when $n \geq 3$, and $\frac{c}{c+1} \geq \frac{2}{3}$ when $c \geq 2$. The final expression is optimized at $c = \frac{2n}{3}$, yielding the desired lower bound. Note that for this choice of c , we indeed have $\delta < 1/2$ when $T \geq 2$.

D.1.2 Proof of Theorem 14

Let the adversarial strategies $\{\mathcal{A}^i\}$ be defined as follows: For each strategy, the adversary independently designates every vertex to be a source, sink, or neither. The adversary then opens directed edges from all source vertices to all sink vertices. For \mathcal{A}^i , with $1 \leq i \leq n$, the adversary designates vertex i to be a source vertex with probability $\frac{c}{n}$, and all other vertices to be source vertices with probability $\frac{c}{n}(1-\delta)$. All vertices are designated to be sink vertices with probability $\frac{d}{n}$. Finally, for \mathcal{A}^0 , the adversary designates all vertices to be source vertices with probability $\frac{c}{n}(1-\delta)$, and sink vertices with probability $\frac{d}{n}$. Successive actions of the adversary are i.i.d. across time steps.

We now derive the following lemmas, which will be used in Proposition 5:

Lemma 13. *For any $i \neq j$ and $1 \leq t \leq T$, we have*

$$\mathbb{E}_i[X_{i,t} - X_{j,t}] = \frac{(n-1)cd}{n^3} \delta.$$

Proof. We compute the expectation of each term separately. Let \mathcal{B}_t and \mathcal{C}_t denote the source and

sink vertices at time t , respectively. Note that $X_{i,t} = \frac{1}{n}$ if $i \notin \mathcal{B}_t$; otherwise, $X_{i,t} = \frac{1+|\mathcal{C}_t|}{n}$. Hence,

$$\begin{aligned}\mathbb{E}_i[X_{i,t}] &= \mathbb{E} \left[\frac{1}{n} \mathbf{1}_{i \notin \mathcal{B}_t} + \frac{1+|\mathcal{C}_t|}{n} \mathbf{1}_{i \in \mathcal{B}_t} \right] \\ &= \frac{1}{n} \left(1 - \frac{c}{n} \right) + \frac{1+(n-1)\frac{d}{n}}{n} \left(\frac{c}{n} \right) \\ &= \frac{1}{n} + \frac{(n-1)cd}{n^3}.\end{aligned}$$

The computation for $X_{j,t}$ is similar:

$$\begin{aligned}\mathbb{E}_i[X_{j,t}] &= \mathbb{E} \left[\frac{1}{n} \mathbf{1}_{j \notin \mathcal{B}_t} + \frac{1+|\mathcal{C}_t|}{n} \mathbf{1}_{j \in \mathcal{B}_t} \right] \\ &= \frac{1}{n} \left(1 - \frac{c}{n}(1-\delta) \right) + \frac{1+(n-1)\frac{d}{n}}{n} \frac{c}{n}(1-\delta) \\ &= \frac{1}{n} + \frac{(n-1)cd}{n^3}(1-\delta).\end{aligned}$$

Taking the difference between these expectations proves the lemma. \square

Lemma 14. *Let $\mathcal{S} \in \mathcal{P}_d$ be a deterministic player strategy, and let $T_i = |\{t : \mathcal{S}_t = \{i\}\}|$. Then we have the upper bound*

$$\sum_{i=1}^n KL(\mathbb{P}_0, \mathbb{P}_i) \leq \frac{c(n-d)}{n(n-c-d)} T \delta^2.$$

Essentially, the Kullback-Leibler divergence is of order $\frac{1}{n}$, because playing a suboptimal vertex provides no information about which vertex is optimal. This is unlike the case of the undirected graph, where the optimal vertex is always more likely to be contained in the feedback that the player receives, and the KL divergence does not decay with n . The proof of Lemma 14 is provided in Appendix D.2.2.

By Proposition 5, we then have

$$\inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}) \geq \frac{(n-1)cd}{n^3} \delta T \left(\frac{n-1}{n} - \delta \sqrt{\frac{T}{2n}} \sqrt{\frac{c(n-d)}{n(n-c-d)}} \right).$$

Finally, we optimize over δ , c , and d . We take

$$\delta = \frac{1}{2} \left(\frac{n-1}{n} \right) \sqrt{\frac{2n}{T}} \sqrt{\frac{n(n-c-d)}{c(n-d)}},$$

to obtain

$$\begin{aligned}\inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}) &\geq \frac{(n-1)cd}{4n^3} \left(\frac{n-1}{n} \right)^2 T \sqrt{\frac{2n}{T}} \sqrt{\frac{n(n-c-d)}{c(n-d)}} \\ &= \frac{1}{2\sqrt{2}} \sqrt{nT} \left(\frac{n-1}{n} \right)^3 \frac{cd}{n^2} \sqrt{\frac{(1-c/n-d/n)}{(c/n)(1-d/n)}} \\ &\geq \frac{1}{16\sqrt{2}} \sqrt{nT} \frac{cd}{n^2} \sqrt{\frac{(1-c/n-d/n)}{(c/n)(1-d/n)}},\end{aligned}$$

where the last inequality uses the bound $\frac{n-1}{n} \geq \frac{1}{2}$. Finally, using the fact that the function

$$f(x, y) = xy \sqrt{\frac{1-x-y}{x(1-y)}}$$

achieves its maximum value of $\frac{1}{3\sqrt{3}}$ when $(x, y) = (\frac{1}{6}, \frac{2}{3})$, we obtain the bound

$$\inf_{\mathcal{S} \in \mathcal{P}} \sup_{\mathcal{A} \in \mathcal{A}} \bar{R}_T(\mathcal{A}, \mathcal{S}) \geq \frac{1}{48\sqrt{6}} \sqrt{nT},$$

when $c = \frac{n}{6}$ and $d = \frac{2n}{3}$.

D.2 Proofs of KL Bounds

In this Appendix, we derive the required upper bounds on the KL divergence between adversarial strategies. We begin by proving a useful technical lemma.

Recall that \mathbb{P}_i denotes the distribution of the edge feedback \mathcal{J}^T under strategy \mathcal{A}^i , and $\mathcal{S} \in \mathcal{P}_d$ is a fixed deterministic player strategy. Also recall that $T_i = |\{t : \mathcal{S}_t = \{i\}\}|$ denotes the number of times vertex i is chosen by the player.

Let \mathbb{P}_i^t denote the distribution of the edge feedback \mathcal{J}^t under strategy \mathcal{A}^i , so $\mathbb{P}_i = \mathbb{P}_i^T$. For each pair of nodes i and v and any $1 \leq t \leq T$, define the function $KL_i^t(v)$ to be the KL divergence between the edge feedback, conditioned on any \mathcal{J}^{t-1} such that $\mathcal{S}_t = \{v\}$:

$$KL_i^t(v) = KL(\mathbb{P}_0^t\{\cdot | \mathcal{J}^{t-1}\}, \mathbb{P}_i^t\{\cdot | \mathcal{J}^{t-1}\}).$$

Note that $KL_i^t(v)$ is indeed a well-defined function of v , since conditioned on \mathcal{J}^{t-1} , the player's action \mathcal{S}_t is deterministic. Hence, the randomness in \mathcal{J}^t is purely due to the stochastic action of the adversary at time t .

Lemma 15. *If $KL_i^t(v)$ is independent of t , we have*

$$KL(\mathbb{P}_0, \mathbb{P}_i) = KL_i(i)\mathbb{E}_0[T_i] + \sum_{j \neq i} KL_i(j)\mathbb{E}_0[T_j]. \quad (51)$$

If in addition $KL_i(i)$ is independent of i , for $1 \leq i \leq n$, and $KL_i(j)$ is constant for all nonzero pairs $i \neq j$, we have

$$\sum_{i=1}^n KL(\mathbb{P}_0, \mathbb{P}_i) = KL_i(i)T + KL_i(j)(n-1)T. \quad (52)$$

Proof. Note that equation (52) follows immediately from equation (51) by summing over i and using the fact that $\sum_{i=1}^n \mathbb{E}_0[T_i] = T$.

To derive equation (51), we use the chain rule for KL divergence:

$$\begin{aligned}
KL(\mathbb{P}_0, \mathbb{P}_i) &= \sum_{t=1}^T \sum_{\mathcal{J}^{t-1}} \mathbb{P}_0\{\mathcal{J}^{t-1}\} KL(\mathbb{P}_0^t\{\cdot|\mathcal{J}^{t-1}\}, \mathbb{P}_i^t\{\cdot|\mathcal{J}^{t-1}\}) \\
&= \sum_{t=1}^T \sum_{v=1}^n \sum_{\mathcal{J}^{t-1}: \mathcal{S}_t=\{v\}} \mathbb{P}_0\{\mathcal{J}^{t-1}\} KL_i^t(v) \\
&\stackrel{(a)}{=} \sum_{t=1}^T \sum_{v=1}^n \mathbb{P}_0\{\mathcal{S}_t = \{v\}\} KL_i(v) \\
&= \sum_{t=1}^T \mathbb{P}_0\{\mathcal{S}_t = \{i\}\} KL_i(i) + \sum_{t=1}^T \sum_{j \neq i} \mathbb{P}_0\{\mathcal{S}_t = \{j\}\} KL_i(j),
\end{aligned}$$

using the assumption that $KL_i^t(v)$ is independent of t in the equation (a). Now we simply recognize that

$$\mathbb{E}_0[T_i] = \mathbb{E}_0 \left[\sum_{t=1}^T \mathbf{1}_{\mathcal{S}_t=\{i\}} \right] = \sum_{t=1}^T \mathbb{P}_0\{\mathcal{S}_t = \{i\}\}$$

to obtain the desired equality. \square

D.2.1 Proof of Lemma 12

Note that $KL_i^t(v)$ is independent of t , since the adversary's actions are i.i.d. across time steps. Furthermore, $KL_i(i)$ is clearly independent of i and $KL_i(j)$ is constant for all pairs $i \neq j$, so equation (52) of Lemma 15 holds.

We first compute an upper bound for $KL_i(i)$. Let X denote the size of the connected component containing i on a particular time step, based on the edges played by the adversary. Then

$$KL_i(i) = KL(\mathbb{P}_0(X), \mathbb{P}_i(X)),$$

where we abuse notation slightly and write $\mathbb{P}_i(X)$ to denote the distribution of X under adversarial strategy \mathcal{A}^i . Also let Y be the indicator variable that i is in the clique selected by the adversary. By the chain rule for the KL divergence,

$$KL(\mathbb{P}_0(X), \mathbb{P}_i(X)) \leq KL(\mathbb{P}_0(X, Y), \mathbb{P}_i(X, Y)).$$

We will derive an upper bound for the latter quantity. In particular, the range of (X, Y) is

$$\{(1, 0), (1, 1)\} \cup \{(m, 1) : 2 \leq m \leq n\}.$$

This leads to the following expression for $KL(\mathbb{P}_0(X, Y), \mathbb{P}_i(X, Y))$:

$$\begin{aligned}
& \left(1 - \frac{c}{n}(1 - \delta)\right) \log \left(\frac{1 - \frac{c}{n}(1 - \delta)}{1 - \frac{c}{n}}\right) \\
& + \left(\frac{c}{n}(1 - \delta) \left(1 - \frac{c}{n}(1 - \delta)\right)^{n-1}\right) \log \left(\frac{\frac{c}{n}(1 - \delta) \left(1 - \frac{c}{n}(1 - \delta)\right)^{n-1}}{\frac{c}{n} \left(1 - \frac{c}{n}(1 - \delta)\right)^{n-1}}\right) \\
& + \sum_{m=2}^n \binom{n-1}{m-1} \frac{c}{n}(1 - \delta) \left(\frac{c}{n}(1 - \delta)\right)^{m-1} \left(1 - \frac{c}{n}(1 - \delta)\right)^{n-m} \\
& \quad \times \log \left(\frac{\frac{c}{n}(1 - \delta) \left(\frac{c}{n}(1 - \delta)\right)^{m-1} \left(1 - \frac{c}{n}(1 - \delta)\right)^{n-m}}{\frac{c}{n} \left(\frac{c}{n}(1 - \delta)\right)^{m-1} \left(1 - \frac{c}{n}(1 - \delta)\right)^{n-m}}\right) \\
& = \left(1 - \frac{c}{n}(1 - \delta)\right) \log \left(\frac{1 - \frac{c}{n}(1 - \delta)}{1 - \frac{c}{n}}\right) \\
& \quad + \sum_{m=1}^n \binom{n-1}{m-1} \frac{c}{n}(1 - \delta) \left(\frac{c}{n}(1 - \delta)\right)^{m-1} \left(1 - \frac{c}{n}(1 - \delta)\right)^{n-m} \log(1 - \delta) \\
& = \left(1 - \frac{c}{n}(1 - \delta)\right) \log \left(\frac{1 - \frac{c}{n}(1 - \delta)}{1 - \frac{c}{n}}\right) + \frac{c}{n}(1 - \delta) \log(1 - \delta).
\end{aligned}$$

Applying the inequality $\log(1 + x) \leq x$ twice, we then obtain

$$\begin{aligned}
KL(\mathbb{P}_0(X), \mathbb{P}_i(X)) & \leq \left(1 - \frac{c}{n}(1 - \delta)\right) \frac{\frac{c\delta}{n}}{1 - \frac{c}{n}} - \frac{c}{n}(1 - \delta)\delta \\
& = \frac{c\delta}{n} \left(\frac{n - c(1 - \delta)}{n - c} - (1 - \delta)\right) \\
& = \frac{c\delta^2}{n - c}.
\end{aligned} \tag{53}$$

The computation for $KL_i(j)$ is similar. Let X denote the size of the connected component containing j , and let \mathcal{C} denote the clique chosen by the adversary. Define the random variable

$$Y = \begin{cases} 0, & \text{if } j \notin \mathcal{C} \\ 1, & \text{if } j \in \mathcal{C} \text{ and } i \notin \mathcal{C} \\ 2, & \text{if } i, j \in \mathcal{C}. \end{cases}$$

Again, it suffices to obtain a bound on $KL(\mathbb{P}_0(X, Y), \mathbb{P}_i(X, Y))$. The range of (X, Y) is

$$\{(1, 0), (1, 1)\} \cup \{(m, 1) : 2 \leq m \leq n - 1\} \cup \{(m, 2) : 2 \leq m \leq n\}.$$

Further note that $\mathbb{P}_0(1, 0) = \mathbb{P}_i(1, 0)$, so we may ignore this term when computing the KL divergence.

We then have following expression for $KL(\mathbb{P}_0(X, Y), \mathbb{P}_i(X, Y))$:

$$\begin{aligned}
& \left(\frac{c}{n}(1-\delta)\right) \left(1 - \frac{c}{n}(1-\delta)\right)^{n-1} \log \left(\frac{\left(\frac{c}{n}(1-\delta)\right) \left(1 - \frac{c}{n}(1-\delta)\right)^{n-1}}{\left(\frac{c}{n}(1-\delta)\right) \left(1 - \frac{c}{n}\right) \left(1 - \frac{c}{n}(1-\delta)\right)^{n-2}} \right) \\
& + \sum_{m=2}^{n-1} \binom{n-2}{m-1} \frac{c}{n}(1-\delta) \left(1 - \frac{c}{n}(1-\delta)\right) \left(\frac{c}{n}(1-\delta)\right)^{m-1} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m-1} \\
& \quad \times \log \left(\frac{\left(\frac{c}{n}(1-\delta)\right) \left(1 - \frac{c}{n}(1-\delta)\right) \left(\frac{c}{n}(1-\delta)\right)^{m-1} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m-1}}{\frac{c}{n}(1-\delta) \left(1 - \frac{c}{n}\right) \left(\frac{c}{n}(1-\delta)\right)^{m-1} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m-1}} \right) \\
& + \sum_{m=2}^n \binom{n-2}{m-2} \left(\frac{c}{n}(1-\delta)\right)^2 \left(\frac{c}{n}(1-\delta)\right)^{m-2} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m} \\
& \quad \times \log \left(\frac{\left(\frac{c}{n}(1-\delta)\right)^2 \left(\frac{c}{n}(1-\delta)\right)^{m-2} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m}}{\left(\frac{c}{n}(1-\delta)\right) \left(\frac{c}{n}\right) \left(\frac{c}{n}(1-\delta)\right)^{m-2} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m}} \right) \\
& = \sum_{m=1}^{n-1} \binom{n-2}{m-1} \frac{c}{n}(1-\delta) \left(1 - \frac{c}{n}(1-\delta)\right) \left(\frac{c}{n}(1-\delta)\right)^{m-1} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m-1} \\
& \quad \times \log \left(\frac{\left(\frac{c}{n}(1-\delta)\right) \left(1 - \frac{c}{n}(1-\delta)\right) \left(\frac{c}{n}(1-\delta)\right)^{m-1} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m-1}}{\frac{c}{n}(1-\delta) \left(1 - \frac{c}{n}\right) \left(\frac{c}{n}(1-\delta)\right)^{m-1} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m-1}} \right) \\
& + \sum_{m=2}^n \binom{n-2}{m-2} \left(\frac{c}{n}(1-\delta)\right)^2 \left(\frac{c}{n}(1-\delta)\right)^{m-2} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m} \\
& \quad \times \log \left(\frac{\left(\frac{c}{n}(1-\delta)\right)^2 \left(\frac{c}{n}(1-\delta)\right)^{m-2} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m}}{\left(\frac{c}{n}(1-\delta)\right) \left(\frac{c}{n}\right) \left(\frac{c}{n}(1-\delta)\right)^{m-2} \left(1 - \frac{c}{n}(1-\delta)\right)^{n-m}} \right) \\
& = \frac{c}{n}(1-\delta) \left(1 - \frac{c}{n}(1-\delta)\right) \log \left(\frac{1 - \frac{c}{n}(1-\delta)}{1 - \frac{c}{n}} \right) + \left(\frac{c}{n}(1-\delta)\right)^2 \log(1-\delta).
\end{aligned}$$

We once again use the inequality $\log(1+x) \leq x$ to obtain

$$\begin{aligned}
KL(\mathbb{P}_0(X), \mathbb{P}_i(X)) & \leq \frac{c}{n}(1-\delta) \left(1 - \frac{c}{n}(1-\delta)\right) \left(\frac{\frac{c\delta}{n}}{1 - \frac{c}{n}} \right) - \left(\frac{c}{n}(1-\delta)\right)^2 \delta \\
& = \left(\frac{c}{n}\right)^2 (1-\delta) \delta \left(\frac{n - c(1-\delta)}{n - c} - (1-\delta) \right) \\
& = \left(\frac{c}{n}\right)^2 (1-\delta) \delta^2 \frac{n}{n - c}.
\end{aligned} \tag{54}$$

Combining inequalities (53) and (54) with equation (52) of Lemma 15, we obtain the bound

$$\begin{aligned}
\sum_{i=1}^n KL(\mathbb{P}_0, \mathbb{P}_i) & \leq \frac{c}{n-c} \delta^2 T + \left(\frac{c}{n}\right)^2 (1-\delta) \delta^2 \frac{n}{n-c} (n-1) T \\
& \leq \frac{c(c+1)}{n-c} T \delta^2,
\end{aligned}$$

completing the proof.

D.2.2 Proof of Lemma 14

Note that $KL_i^t(v)$ is independent of t , since the adversary's actions are i.i.d. across time steps. Furthermore, $KL_i(i)$ is clearly independent of i and $KL_i(j)$ is constant for all pairs $i \neq j$, so equation (52) of Lemma 15 holds.

Note that when $\mathcal{S}_t = \{j\}$, the distribution of the feedback \mathcal{J}_t is the same under $\mathbb{P}_0^t\{\cdot|\mathcal{J}^{t-1}\}$ and $\mathbb{P}_i^t\{\cdot|\mathcal{J}^{t-1}\}$, since the vertex i is chosen to be a sink vertex with the same probability $\frac{d}{n}$ under both \mathcal{A}^0 and \mathcal{A}^i . Hence, $KL_i(j) = 0$.

To compute $KL_i(i)$, let X denote the size of the infected component containing i when $\mathcal{S}_t = \{i\}$, and define the random variable

$$Y = \begin{cases} 0, & \text{if } i \text{ is a sink vertex} \\ 1, & \text{if } i \text{ is a source vertex} \\ 2, & \text{otherwise.} \end{cases}$$

As in the proof of Lemma 12, we will upper-bound $KL(\mathbb{P}_0(X, Y), \mathbb{P}_i(X, Y))$, leading to an upper bound on $KL(\mathbb{P}_0(X), \mathbb{P}_i(X))$. The range of (X, Y) is

$$\{(1, 0), (1, 2)\} \cup \{(m, 1) : 2 \leq m \leq n\}.$$

We then have the following expression for $KL(\mathbb{P}_0(X, Y), \mathbb{P}_i(X, Y))$:

$$\begin{aligned} & \frac{d}{n} \log \left(\frac{d/n}{d/n} \right) + \left(1 - \frac{c}{n}(1 - \delta) - \frac{d}{n} \right) \log \left(\frac{1 - \frac{c}{n}(1 - \delta) - \frac{d}{n}}{1 - \frac{c}{n} - \frac{d}{n}} \right) \\ & + \sum_{m=2}^n \binom{n-1}{m-1} \frac{c}{n}(1 - \delta) \left(\frac{d}{n} \right)^{m-1} \left(1 - \frac{d}{n} \right)^{n-m} \\ & \quad \times \log \left(\frac{\frac{c}{n}(1 - \delta) \left(\frac{d}{n} \right)^{m-1} \left(1 - \frac{d}{n} \right)^{n-m}}{\frac{c}{n} \left(\frac{d}{n} \right)^{m-1} \left(1 - \frac{d}{n} \right)^{n-m}} \right), \\ & = \frac{n - c - d + c\delta}{n} \log \left(\frac{n - c - d + c\delta}{n - c - d} \right) \\ & \quad + \frac{c}{n}(1 - \delta) \log(1 - \delta) \sum_{m=2}^n \binom{n-1}{m-1} \left(\frac{d}{n} \right)^{m-1} \left(1 - \frac{d}{n} \right)^{n-m} \\ & \leq \frac{n - c - d + c\delta}{n} \log \left(\frac{n - c - d + c\delta}{n - c - d} \right) + \frac{c}{n}(1 - \delta) \log(1 - \delta). \end{aligned}$$

Using the inequality $\log(1 + x) \leq x$, we then have

$$KL(\mathbb{P}_0(X), \mathbb{P}_i(X)) \leq \frac{c\delta(n - c - d + c\delta)}{n(n - c - d)} - \frac{c}{n}(1 - \delta)\delta = \frac{c(n - d)}{n(n - c - d)}\delta^2.$$

Applying Lemma 15 completes the proof.

References

L. A. Adamic and E. Adar. Friends and neighbors on the Web. *Social Networks*, 25(3):211 – 230, 2003. ISSN 0378-8733.

- A. Ageev, R. Hassin, and M. Sviridenko. A 0.5-approximation algorithm for MAX DICUT with given sizes of parts. *SIAM Journal on Discrete Mathematics*, 14(2):246–255, 2001.
- R. Angell and G. Schoenebeck. Don’t be greedy: Leveraging community structure to find high quality seed sets for influence maximization. *arXiv preprint arXiv:1609.06520*, 2016.
- C. Borgs, M. Brautbar, J. Chayes, and B. Lucier. Maximizing social influence in nearly optimal time. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 946–957. SIAM, 2014.
- S. Bubeck and N. Cesa-Bianchi. Regret analysis of stochastic and nonstochastic multi-armed bandits. *Foundations and Trends in Machine Learning*, 5(1):1–122, 2012.
- N. Cesa-Bianchi and G. Lugosi. *Prediction, Learning, and Games*. Cambridge University Press, New York, NY, 2006.
- W. Chen, Y. Wang, and S. Yang. Efficient influence maximization in social networks. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 199–208. ACM, 2009.
- W. Chen, C. Wang, and Y. Wang. Scalable influence maximization for prevalent viral marketing in large-scale social networks. In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 1029–1038. ACM, 2010a.
- W. Chen, Y. Yuan, and L. Zhang. Scalable influence maximization in social networks under the linear threshold model. In *Proceedings of the 2010 IEEE International Conference on Data Mining, ICDM ’10*, pages 88–97, Washington, DC, USA, 2010b. IEEE Computer Society.
- W. Chen, T. Lin, Z. Tan, M. Zhao, and X. Zhou. Robust influence maximization. *arXiv preprint arXiv:1601.06551*, 2016a.
- W. Chen, Y. Wang, Y. Yuan, and Q. Wang. Combinatorial multi-armed bandit and its extension to probabilistically triggered arms. *Journal of Machine Learning Research*, 17(50):1–33, 2016b.
- J. Cheng, L. Adamic, P. A. Dow, J. Kleinberg, and J. Leskovec. Can cascades be predicted? In *Proceedings of the 23rd International Conference on WWW*, pages 925–936. ACM, 2014.
- P. Domingos and M. Richardson. Mining the network value of customers. In *Proceedings of the seventh ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 57–66. ACM, 2001.
- R. Durrett. *Lecture Notes on Particle Systems and Percolation*. Wadsworth & Brooks/Cole statistics/probability series. Wadsworth & Brooks/Cole Advanced Books & Software, 1988.
- U. Feige, V. S. Mirrokni, and J. Vondrak. Maximizing non-monotone submodular functions. *SIAM Journal on Computing*, 40(4):1133–1153, 2011.
- M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.
- A. Gupta, A. Roth, G. Schoenebeck, and K. Talwar. Constrained non-monotone submodular maximization: Offline and secretary algorithms. In *Internet and Network Economics*, pages 246–257. Springer, 2010.

- X. He and D. Kempe. Robust influence maximization. *arXiv preprint arXiv:1602.05240*, 2016.
- M. Hecker, S. Lambeck, S. Toepfer, E. Van Someren, and R. Guthke. Gene regulatory network inference: data integration in dynamic models—a review. *Biosystems*, 96(1):86–103, 2009.
- D. Kempe, J. Kleinberg, and É. Tardos. Maximizing the spread of influence through a social network. In *Proceedings of the Ninth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '03, pages 137–146, New York, NY, USA, 2003. ACM. ISBN 1-58113-737-0.
- D. Kempe, J. Kleinberg, and É. Tardos. Influential nodes in a diffusion model for social networks. In *Automata, languages and programming*, pages 1127–1138. Springer, 2005.
- W. O. Kermack and A. G. McKendrick. A contribution to the mathematical theory of epidemics. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 115(772):700–721, 1927.
- J. Khim, V. Jog, and P. Loh. Computing and maximizing influence in linear threshold and triggering models. In *Advances in Neural Information Processing Systems*, 2016.
- E. J. Lee, S. Kamath, E. Abbe, and S. R. Kulkarni. Spectral bounds for independent cascade model with sensitive edges. In *2016 Annual Conference on Information Science and Systems (CISS)*, pages 649–653, March 2016.
- S. Lei, S. Maniu, L. Mo, R. Cheng, and P. Senellart. Online influence maximization. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '15, pages 645–654, New York, NY, USA, 2015. ACM. ISBN 978-1-4503-3664-2.
- R. Lemonnier, K. Scaman, and N. Vayatis. Tight bounds for influence in diffusion networks and application to bond percolation and epidemiology. In *Advances in Neural Information Processing Systems*, pages 846–854, 2014.
- R. Lemonnier, K. Scaman, and N. Vayatis. Spectral Bounds in Random Graphs Applied to Spreading Phenomena and Percolation. *ArXiv e-prints*, March 2016.
- J. Leskovec, L. A. Adamic, and B. A. Huberman. The dynamics of viral marketing. *ACM Transactions on the Web (TWEB)*, 1(1):5, 2007.
- D. Liben-Nowell and J. Kleinberg. The link-prediction problem for social networks. *Journal of the American Society for Information Science and Technology*, 58(7):1019–1031, 2007. ISSN 1532-2890. doi: 10.1002/asi.20591. URL <http://dx.doi.org/10.1002/asi.20591>.
- T. Liggett. *Interacting Particle Systems*, volume 276. Springer Science & Business Media, 2012.
- N. Movarraei and S. Boxwala. On the number of paths of length 5 in a graph. *International Journal of Applied Mathematical Research*, 4(1):30–51, 2015a.
- N. Movarraei and S. Boxwala. On the number of paths of length 6 in a graph. *International Journal of Applied Mathematical Research*, 4(2):267–280, 2015b.
- N. Movarraei and M. Shikare. On the number of paths of lengths 3 and 4 in a graph. *International Journal of Applied Mathematical Research*, 3(2):178–189, 2014.
- H. Narasimhan, D. C. Parkes, and Y. Singer. Learnability of influence in networks. In *Advances in Neural Information Processing Systems*, pages 3186–3194, 2015.

- G. L. Nemhauser and L. A. Wolsey. Best algorithms for approximating the maximum of a submodular set function. *Mathematics of operations research*, 3(3):177–188, 1978.
- M. E. J. Newman. The structure and function of complex networks. *SIAM review*, 45(2):167–256, 2003.
- K. Scaman, R. Lemonnier, and N. Vayatis. Anytime influence bounds and the explosive behavior of continuous-time diffusion networks. In *Advances in Neural Information Processing Systems*, pages 2017–2025, 2015.
- O. Sporns. The human connectome: A complex network. *Annals of the New York Academy of Sciences*, 1224(1):109–125, 2011.
- M. Streeter and D. Golovin. An online algorithm for maximizing submodular functions. *Technical report*, pages 1–35, 2007.
- S. Vaswani, L. Lakshmanan, and M. Schmidt. Influence maximization with bandits. *arXiv preprint arXiv:1503.00024*, 2015.
- Z. Wen, B. Kveton, and M. Valko. Influence maximization with semi-bandit feedback. *CoRR*, abs/1605.06593, 2016. URL <http://arxiv.org/abs/1605.06593>.
- C. Zhou, P. Zhang, J. Guo, and L. Guo. An upper bound based greedy algorithm for mining top- k influential nodes in social networks. In *Proceedings of the 23rd International Conference on World Wide Web*, pages 421–422. ACM, 2014.